THÈSE

En vue de l'obtention du

DOCTORAT DE L'UNIVERSITÉ DE TOULOUSE

Délivré par l' Université Toulouse 1 Capitole
Discipline : Sciences Economiques

Présentée et soutenue par

Nicolas DUPUIS

Le 27 avril 2015

Titre :

Essays in Industrial Organization

Ecole doctorale : Toulouse School of Economics

Unité de recherche : GREMAQ – TSE

Directrice de Thèse : Marc IVALDI

JURY

Alessandro GAVAZZA, professeur, London School of Economics
Marc IVALDI, professeur, Université Toulouse 1
Bruno JULLIEN, professeur, Université Toulouse 1
Peter Norman SØRENSEN, professeur, University of Copenhagen

Thesis advisor: Professor Marc Ivaldi Nicolas Dupuis

Essays in Industrial Organization

Abstract

In this dissertation, I investigate several questions of interest in the field of Industrial Organization.

My contribution to the literature on cheap talk is both theoretical and empirical. By introducing endogeneity in the precision with which the audience observes the state of the world, I am able to identify two new effects distorting the experts' incentives to reveal their private information: a barrier to entry effect and a guru effect. Moreover, I develop an innovative way to estimate the experts' ability and biases in the context of a reputational cheap-talk game.

I contribute to the literature on switching costs and tacit collusion by investigating the impact of consumers characteristics on the feasibility of tacit collusion in industries with switching costs. The sophistication and lifespan of consumers have a strong impact the lower and upper bounds of the switching costs and discount factors which can sustain tacit collusion.

Finally, I contribute to the literature on revenue management and intertemporal price discrimination by constructing a realistic and rich model of revenue management and studying the welfare impact of the pricing strategy. The model accommodates heterogeneous consumers and competition between revenue managers. The practice generally increases both the profits of the company and the consumer surplus as it allows the company to sell more items.

In my first chapter, *Avoiding Judgement by Recommending Inaction: Beliefs Manipulation and Reputational Concerns*, co-authored with Fanny Camara, we investigate experts' incentives to reveal their private information on the state of the world when their audience observes the state of the world with some

endogenous noise. Experts are concerned by the audience's perception about their ability to get private information that is informative about the true state of the world. The perception of the audience changes according to whether or not an expert revealed a piece of information corresponding to the state of the world. However the audience cannot always observe perfectly the state of the world and the precision with which it observes the state depends on the ex-ante belief of the audience. To take an example, if everyone expects a movie to be great, it receives a large number of spectators who can then share their opinions about it. In the end, everyone gets a fairly precise idea on whether or not the movie was good. If everyone expects a movie to be dull, no one sees it but no one can tell whether or not the movie was actually dull.

The incentives of the expert to report her true opinion are distorted in two new ways. First, as she is unsure about the accuracy of her information, she has greater incentives to comfort her audience in its opinion that the movie is dull. By doing so, she prevents her audience to update unfavourably its perception about her ability. We call this effect the barrier to entry effect. Second, she has greater opportunities to discourage her audience to see the movie if her audience trusts her. The audience of the expert trusts her when it thinks she has high quality information, i.e. when the expert has a high perceived ability. The expert's incentives to manipulate the audience's beliefs increase with her perceived ability. But, as the audience correctly anticipates the experts strategy in equilibrium, experts with high perceived abilities lose their credibility. We call this effect the guru effect. Competition between experts alleviates the guru effect and can increase the amount of valuable information transmitted to the audience. Our model has applications in markets for experience goods but also to evaluate investments advices.

In my second chapter, *Structural Estimation of Expert Strategic Bias: the Case of Movie Reviewers*, co-authored with Fanny Camara, we use the structure of our theoretical model to estimate the ability and biases of movie reviewers. This empirical work develops the first structural approach to estimate

reputational cheap-talk games and applies it to movie reviews released in the US between 2004 and 2013.

We jointly identify prior beliefs on the quality of movies and the reviewers' biases and abilities. For given biases and abilities, higher prior movies tend to get more unanimously positive reviews. Once we identify the prior, we retrieve both abilities and biases using a partition on the prior space: priors either fall within or outside a truthful revelation set. The truthful revelation set is the subset of the prior space on which any expert of any ability truthfully reveals her information in equilibrium. Its position in the prior space also depends on the reputation of the expert. Inside the truthful revelation set, absent of bias, priors and experts' abilities determine the distribution of observed recommendations. We can therefore recover abilities for given priors. Outside the truthful revelation set, for given priors and abilities, experts' biases are directly identified from the observed recommendations. We observe variation in the reputation of each reviewer over time which shifts their truthful revelation set and allows us to identify the effect of reputation on each reviewer's bias. Our estimation strategy directly proceeds from the identification as we use the same partition structure over the set of priors in a maximum likelihood approach.

We find that abilities, defined by the probability for the reviewer to receive an accurate signal on the movie quality, range from 60% to 89%. The bias, defined as the probability for the reviewer to report a recommendation that contradicts her private signal, is 10% on average and can go up to 40%. We find a modest but significant impact of the reviewer's reputation which translates into the reviewer being more reluctant to report truthfully high signals compared to low ones. Our estimation takes into account and quantifies potential conflicts of interest that might arise when the movie reviewer belongs to the same media outlet as the film under review. Out-of-sample predictions confirm that movie reviewers do have reputational concerns.

Our contributions are threefold. We are the first to estimate a cheap talk game. We generalize the identification and estimation approach for voting games developed by Iaryczower and Shum (2012) by allowing experts to have strategic and then endogenous bias. This paper also contributes to the literature

on estimating conflict of interest resulting from the increasing concentration of the media industry.

In my third chapter, *Market Characteristics and Implicit Contracts when Consumers have Switching Costs*, co-authored with Guillem Roig, we examine how the sustainability of tacit collusion changes with the sophistication and lifespan of consumers in a market characterized by switching costs. We consider a dynamic framework in which two competing firms post a unique price per period and use simple grim-trigger strategies. In this framework, collusion might be sustainable for a given level of switching costs and discount factors if consumers are short-sighted but not if they are forward-looking. But this result might not hold depending on whether consumers have the possibility to repeatedly buy the product or not. In the presence of two representative consumers, maintaining a fully collusive outcome is more difficult when consumers are forward-looking. This is not necessarily the case when we consider a continuum of these forward-looking consumers.

We contribute to the literature on switching costs and tacit collusion by solving the game when consumers are forward-looking and firms use simple strategies with deterministic prices. We also highlight the importance of the hypotheses chosen to study the feasibility of tacit collusion in a given market as one factor in the market environment can have a considerable impact on the collusive outcome. Our work could help to determine which markets are more likely to harbour tacit collusion.

In my fourth chapter, *A Welfare Assessment of Revenue Management*, co-authored with Pr. Marc Ivaldi and Pr. Jerome Pouyet, we study the welfare impact of revenue management, i.e. intertemporal price discrimination when the product availability is limited both in time and quantity, and consumers' arrival is random.

We develop a theoretical and rich model of revenue management which allows for heterogeneity in product characteristics, capacity constraints, consumer preferences, and probabilities of arrival. We also

introduce competition between revenue managers. We solve this model computationally and recover the optimal pricing strategies. We find that revenue management is welfare enhancing. Prices change during the booking period as the revenue manager reacts to the past sales. If she has a lot of time to sell her remaining stock, she wants to increase her price in the hope that consumers with a high valuation for her product show up. But she also faces a time limit to sell her products and when this time limit comes close, she wants to decrease her prices to sell her last remaining items. Compared to a fixed price, profits increase as revenue management offers more leeway to the seller. Total consumer surplus also increases for a wide range of specifications, as revenue management raises the number of sales. In the presence of heterogeneous consumers, consumers with low price sensitivity subsidize ones with high price sensitivity when demand is low but both types benefit from the practice when demand is high.

The practice is particularly relevant, and widely spread, in the transport industry. Our realistic framework provides a general answer to the question of the welfare impact of revenue management but is also flexible and can be applied to a given market to conduct specific welfare analyses. We also contribute to the literature by introducing competition between revenue managers.

Contents

Listing of figures

List of Tables

To my Parents.

1

Avoiding Judgement by Recommending Inaction: Beliefs Manipulation and Reputational Concerns

co-authored with Fanny Camara

Abstract

We generalize a reputational cheap-talk model by introducing an endogenous level of receivers' monitoring of the state of the world: when the audience strongly believes that the state is low, the precision with which it observes the state of the world is also low. For instance, the quality of experience goods is revealed to the market through consumer reviews. When potential consumers have a low prior belief on the quality of the good, they are less likely to buy and then review it, generating a noisy ex-post feedback for the audience. In these situations, experts who have low skills or do not know their ability misreport their signals in equilibrium as soon as they have a strong prior on the state of the world. Compared to models with perfect observability of the state of the world, we identify two additional channels through which reports are biased: an incentive to overreport the low state, which induces a lower quality in receivers' monitoring, hence a lower capability to update their perception on the expert's ability; and the imperfect monitoring reinforces the incentive to stick to the receivers' prior belief. Also, an expert with a higher perceived ability is less likely to reveal a high signal when the prior is low. The expert can directly impact the receivers' beliefs about the state of the world through her message, creating what we call a *guru effect*. Contrary to the predictions of reputational cheap-talk games with perfect ex-post observability of the state of the world, the expert's perceived ability and competition among experts affect equilibrium reporting behaviours as they enter in the way experts can manipulate beliefs. Competition between experts reduces the incentives to overreport the low signal and counterbalances the perverse effects of a high perceived ability.

1.1 Introduction

Experts do not necessarily offer their expertise on matters that are perfectly observable to their audience ex-post. Sometimes, the extent to which those matters are observable depends on the advice of the expert. If Sally advises John not to read a book and John trusts Sally, John cannot ascertain the quality of the book. If the advice was to read the book, John would know for sure whether he likes it or not. In such instances, it is not an easy task for the audience to verify the accuracy of the expert's advice. If John cannot ascertain the quality of the book, he cannot tell if Sally's advice was a good one or not. We study experts' incentives to divulge their true opinion in such frameworks.

We develop a reputational cheap-talk game featuring an expert with noisy private information about the state of the world who transmits this private information to a group of receivers. The precision of the expert's private information depends on her ability to be well informed. The expert wishes to maximize the audience's perception on her ability to get accurate signals on the true state of the world, or perceived ability. We contribute to the literature on reputational cheap-talk game by addressing cases in which the state of the world is observed with some endogenous noise. The state of the world predicted by the expert can be the quality of an experience good, the profitability of an investment, or the quality of a scientific article.

The quality of an experience good assessed by a reviewer is revealed to the market through the aggregation of individual consumer experiences, which can take the form of simple word-of-mouth communication or aggregate consumer reviews online. For instance, users of websites aggregating movie reviews such as *rotten tomatoes* can post their own appreciation of a movie. The website then aggregates reviews and presents an audience score next to the average ratings of the critics. The true quality is not observed perfectly: the lower the share of the consumers who actually purchased and experienced the product, the less precise the aggregate signal received by the market about the quality of the product. The profitability of an investment project evaluated by a financial analyst cannot be observed by potential investors unless enough agents invest in the project. The quality of a scientific article evaluated by a journal editor can be revealed to the scientific community only if the editor decides to publish it. In all the above cases, the recommendation made by the expert plays a crucial role in allowing the market to observe a more or less precise signal about the state.

Games of cheap-talk feature several equilibria, including the babbling equilibrium in which the audience does not listen to the expert and the expert does not transmit any credible information to the audience. The babbling equilibrium maximizes the loss of valuable information. Among all the possible equilibria of the game, we are interested in the equilibrium that minimizes the loss of valuable information.

The structure of the most informative equilibrium depends on whether the expert knows her own ability or not, i.e. knows how good she is at receiving signals corresponding to the true state of the world or not.

If the expert knows her own ability and this ability can either be low or high, the high ability expert reveals truthfully her private signals on the whole range of initial prior beliefs as she wants to separate herself from low ability experts. Indeed, a low ability expert would tell the truth on an intermediate range of prior beliefs, i.e. when the probability that the state of the world is not too close to 0 or 1, but would lie with some probability outside when her private signal contradicts the most likely state of the world. For instance, if a high state is very likely and the low ability expert receives a signal indicating a low state, she conceals this signal with some positive probability. We call truthful revelation set the subset of prior beliefs on which experts of both abilities reveal their signals. The width of the truthful revelation set gives an indication on how much information is transmitted in equilibrium.

When the expert does not know her own ability, she shares the same information as the audience about how good she is at receiving precise signals. The most informative equilibrium features truthtelling in an intermediate range of prior beliefs and babbling outside for experts of both abilities, meaning that depending on the prior belief, the message of the expert to her audience is either completely truthful of totally independent from her private information. Again, we call truthful revelation set the subset of prior beliefs on which truthful revelation of the private signal is an equilibrium.

The structure of the equilibrium coincides with the one in reputational cheap-talk models in which the state of the world is perfectly observed ex-post. However, adding the assumption of the endogenous precision of the ex-post aggregate signal changes the shape of the truthful revelation set in two significant ways.

First, revelation of information is more likely for prior beliefs in favour of a high state of the world than for prior beliefs in favour of a low state. Indeed, the truthful revelation set is not symmetric around the average prior but shifted towards high prior beliefs on the high state of the world. As the precision of the ex-post aggregate signal is lower when the audience expects a low state of the world, the expert has the extra incentive to comfort prior beliefs indicative of a low state. By doing so, the expert prevents the audience from updating negatively its perception about her ability. Applied to experience goods, experts' reviews favour well-established brands and reinforce pessimistic beliefs about new entrants, creating a barrier to entry. We call *barrier to entry effect* the asymmetry of the truthful revelation set with respect to prior beliefs. This barrier to entry effect only arises when the expert is able to manipulate demand, i.e. when the precision of the ex-post aggregate signal is endogenous.

Second, for a given true ability, an expert perceived to be more able has less incentives to reveal a positive signal when the state of the world is likely to be low. If the agents perceive the expert as someone receiving extremely precise signals about the state of the world, they have a high confidence in her message. But the trust of the audience creates a perverse incentive for the expert: she can use this trust to her advantage by consistently confirming prior beliefs indicative of a low state of the world. In equilibrium, the audience who anticipates the strategy of the expert does not trust her anymore and truthful revelation does not occur. But it would have, if the expert had a lower perceived ability. We call *guru effect* the negative impact of a high reputation on the truthful revelation of private signals.

In addition to the barrier to entry effect and guru effect, the endogenous precision of the ex-post aggregate signal induces more conformism towards the initial expectations of the audience. Moreover, competition kills the option for each expert to manipulate demand by sending unfavourable recommendations when the prior belief is indicative of a low state of the world. Indeed, as the number of expert increases, so does the probability that at least one of them is of higher ability and finds it optimal to reveal a high signal truthfully. By doing so, the high ability expert increases the accuracy of the ex-post aggregate signal and induces more honesty from lower ability experts.

To ease the exposition of the model and of our results, we use a specific functional form to describe how the precision of the aggregate signal changes with the expectation of the audience: if after observing the message of the expert, the audience belief on a high state of the world falls below a threshold, the ex-post aggregate signal is completely uninformative. In the other case, the precision of the aggregate signal equals one, meaning that the audience observes the true state of the world. When the expert does not know her own ability, we prove our results for a completely general relation between the precision of the aggregate signal and the expectation of the audience. When the expert knows her ability however, we could not find a closed-form solution for the bounds of the truthful revelation set in the more general setting. Fortunately, numerical examples seem to indicate that our results on the barrier to entry effect and guru effect hold in the general case.

We use the specific functional form for the precision of the aggregate signal in the introduction of our model in section 1.3 and in the exposition of our results in sections 1.4 to 1.6. We discuss the generalization of our results and the impact of competition between experts in section 1.7.

1.2 Related Literature

The literature on experts and cheap-talk is wide. Several papers, e.g. Morris [38], examine the case of an expert advocating for the less desirable solution just to signal her type and gain the trust of the principal in the future periods. In our paper, we examine a reduced-form model and are only interested about deviations from truthful revelation.[1]

Our approach is closer to the ones used in Ottaviani and Sørensen [41] and Ottaviani and Sørensen

[1]Although we do not exclude equilibria in which experts with high skills blindly recommend low prior products just to prove their value.

[42], and in Gentzkow and Shapiro [23]. Our model is largely inspired from the binary experiment model in Ottaviani and Sørensen [42]: they consider the case in which the expert has a signal on the state of the world, makes her recommendation, and the true state of the world is ultimately known by the receivers who update the reputation of the expert. We model our game exactly along their lines except that we allow for an endogenous level of receivers' monitoring of the state of the world. In Gentzkow and Shapiro [23], a similar model is applied to the media industry where the expert is a journalist. While they model imperfect observation of the state of the world, they do not endogenize the degree of noise. Introducing the endogeneity allows us to recover the barrier to entry and guru effects and to show that competition does affect the reporting behaviour of the experts.

Holmström [26] studies the case of the investment choices of a manager. The manager can be talented or not and her talent determines her capability to choose successful investments. This model is close to ours because if the investment is not chosen, the quality of the project is not observed. The story however differs from ours as in Holmström [26], the incentives not to invest never depend on the initial perceived ability of the expert which is our guru effect. Such incentives arise because the expert does not manipulate directly the observability of the state of the world.

Our paper is also linked to a growing literature in finance. Hence, Mariano [34] develops a model of rating agencies reporting on the profitability of projects. Although our frameworks are similar, our model is more general: we do not assume that the high type expert is perfectly informed and the low type expert completely uninformed. The precision of the ex-post aggregate signal is also completely endogenous in our case. It depends on the message of the expert and on the initial prior, which implies that the expert's perceived ability affects her strategy. In contrast, in Mariano [34] a bad project fails with certainty in the end and a good project succeeds for sure in the case of a good evaluation but fails with some fixed probability in the case of a bad one. This is clearly problematic in the sense that in equilibrium, the receivers may not believe the expert when she is not truthful, therefore there is no objective reason why the evaluation should negatively affect the outcome of the project in that case. Moreover, if the prior is sufficiently low and the expert not credible enough, the project might not attract investments even in the case of a good evaluation. In that case, even a good project might fail.

1.3 The Model

An expert observes a private signal over the state of the world and sends a recommendation to a pool of agents or audience. The agents observe the state of the world with some noise which allows them to update their belief about the true state. In the process, they also update their perception of the expert's ability: if the agents think the expert has received a signal close to the true state of the world, the expert's perceived ability increases and vice versa.

The precision with which the state is observed increases with the number of receivers who observe an individual signal on the state. The idea is that the combined opinions of these agents form an aggregate signal for the others. For instance, people who read a book can share their opinions about its quality by posting comments on the internet or by talking to their friends. In the first case, the pool consists in all internet users, in the second, it consists in the group of friends. The precision of the aggregate signal increases with the number of agents in the pool who chose to observe a signal on the state of the world, i.e. read the book. We expose in Appendix A.1 a formal argument on how agents' individual opinions aggregate to a single ex-post signal whose precision depends on the number of agents.

We indirectly model the dependence of the precision of the aggregate signal on the choice of agents by considering that the precision increases with the agents' belief about a high state of the world after observing the recommendation of the expert. If people expect the book to be great, a large part of the pool will read it.

We restrict the way the aggregate signal depends on the choice of the agents. If after observing the message of the expert their belief falls under a given threshold, too few choose to observe a signal on the state of the world and the ex-post aggregate signal is completely uninformative. If, on the other hand, their belief falls above the threshold, the number of agents who observe a signal on the state of the world exceeds a critical mass and the ex-post aggregate signal is completely informative. If the expectations on the quality of the book are too low, no one reads the book and reports her opinion about it. If the expectations are high enough, a sufficient number of people reads the book for their combined opinions to be perfectly informative about its quality. [2]

When sending her message, the expert faces three possibilities: a) whatever her message, the agents' belief falls under the threshold; b) whatever her message, the agents' belief falls above the threshold; c) the agents' belief falls under the threshold if the expert sends a message indicating a low state of the world and above if she sends the opposite message.

The two first cases correspond to states of the world with prior beliefs associated to them that are either extremely low or extremely high. In the two first cases, the expert cannot control the precision of the ex-post aggregate signal through her message. When the prior belief is extremely low, agents never observe an informative aggregate signal about the state and the expert cannot be credible. When the prior belief is extremely high, agents always observe the true state in the end and the expert sends a message corresponding to what she thinks is the most likely state of the world. We go back to a classic reputational cheap-talk model. In that situation, the expert might disregard her private signal when it contradicts her

[2] This functional form for the precision of the aggregate signal has the advantage to yield closed-form solutions and clear intuitions. Most of our results extend to a general functional form.

initial belief.

The third case corresponds to states of the world with prior beliefs that are intermediate. With such prior beliefs, the expert can manipulate the precision of the aggregate signal using her message, as long as the audience believes her. This possibility creates a strong incentive to send messages indicating a low state of the world as the expert ensures that the audience cannot update its perception about her ability. Because of this incentive, the expert loses her credibility in equilibrium when the prior is low.

We now proceed to a more formal exposition of the model.

1.3.1 Information Structure

The State of the World

The true state of the world is denoted $\theta \in \{0, 1\}$. θ_0 or $\theta = 0$ means that the state of the world is low. Initially, both the agents and the expert share a common prior on the true state of the world, which we denote $\mu_i = p(\theta_i)$, for $i = 0, 1$.[3] In the course of their interaction with the expert, agents form an updated prior belief, which corresponds to their belief on the state of the world after observing the message of the expert and is denoted $\nu_i = p(\theta_i|message)$, for $i = 0, 1$.

Agents do not necessarily observe the true state of the world with probability 1. Instead, at the end of the game, all players observe an aggregate signal $X \in \{0, 1\}$. $X_i|\theta_i$ follows a binomial distribution with probability of success τ, which is a function of the belief of the audience on the high state of the world after observing the message of the expert.

$$p(X_i|\theta_i) = \tau(\nu_1) \in \left[\frac{1}{2}, 1\right]$$

The All or Nothing Aggregate Signal

The aggregate signal is completely random, hence uninformative, if the agents' belief on the high state of the world falls under a threshold and completely informative otherwise. τ is equal to $\frac{1}{2}$ if the agents' belief on the high state of the world is too low and to 1 if the agents' belief is sufficiently high.

$$\tau(\nu_1) = \frac{1 + 1(\nu_1 \geq \chi^{-1})}{2}$$

and $\chi^{-1} \in [0, 1]$ is the threshold marking the passage from a completely uninformative to a completely informative aggregate signal.

[3] Although technically imprecise, for the sake of clarity we sometimes use the phrase "a high prior on the state of the world" to refer to high prior beliefs on the state of the world being high. Conversely, we use "a low prior on the state of the world" to

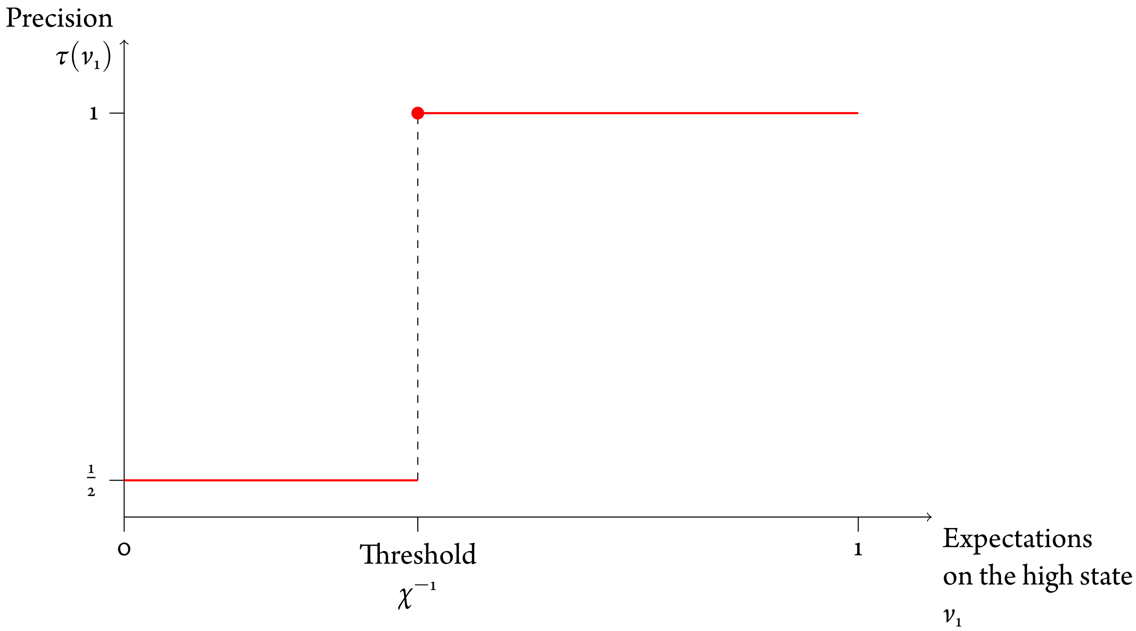

Figure 1.3.1: The All or Nothing τ Function

1.3.2 The Players

The expert receives a private signal $s \in \{0, 1\}$ indicative of the state of the world. The private signal is noisy and a more able expert receives more precise signals. Formally, the ability of the expert, $t \in \{l, h\}$ with $1/2 < l < h < 1$, is the probability with which she gets the correct signal about the state: $t = p(s_i|\theta_i) \quad \forall i \in \{0, 1\}$. As the lowest possible ability l is strictly greater than $\frac{1}{2}$, a signal is always partially informative. When she receives s_0, the expert update her belief on the state of the world and allocate more weight on the low state than before.

We denote $\hat{t}$ the expected ability of the expert with respect to the expert's information set. The agents do not know the ability of the expert but form a belief on her expected ability at the beginning of the game: $\tilde{t} = hp(h) + l(1 - p(h))$. We call this belief perceived ability. If $\hat{t} = t$, the expert knows her own ability. If $\hat{t} = \tilde{t}$, both expert and receivers share the same information, i.e. the expert only knows her perceived ability.

Once she has received her private signal, the expert updates her belief on the state of the world and gets her private posterior $p(\theta|s, \hat{t})$.

We assume that the expert receives zero utility when perceived to have a low ability and a utility equal to 1 when perceived to have a high ability. The goal of the expert is to maximize her expected utility, hence

refer to high prior beliefs on the state of the world being low.

her perceived ability. As a high ability expert receives private signals reflecting the true state of the world more regularly, the audience's perception about her ability increases if the pool of agents believes the expert has received a signal reflecting the true state of the world.

In our setting, as in Ottaviani and Sørensen [42], the group of agents infers the signal received by the expert using her message and a guess on her strategy. They can then update their perception on the ability of the expert. Although both the pool of agents and the expert initially share the same information about the state of the world, i.e. their common prior belief, as soon as the expert receives her private signal, beliefs of the expert and the agents can differ. To avoid confusion, we add the superscripts a for the audience or agents and e for the expert whenever we consider it necessary.

Timing of the game

Figure 1.3.2 presents the timing of the game. Common features of our model and the classical reputational cheap talk model are the setup and expert phases. Part (a) of the figure displays the timing specific to our model. We present the timing of the game with perfect observability of the state of the world, or benchmark model, in part (b). The benchmark model is a special case of our model where the precision of the aggregate signal X is 1. (Which we can model by choosing $\chi \to \infty$.)

The steps of the game are:

Step 1 The state of the world is drawn. Both the expert and the pool of agents share a common prior on its value: $p(\theta_i)$. The perceived ability of the expert is $\tilde{t}$. The expert's information about her ability is given by $\hat{t}$.

Step 2 The expert observes privately her signal s. She gets her private posterior belief on the state of the world $p^e(\theta_i|s,\hat{t})$.

Her message $m \in \{0,1\}$ to the agents is supposed to convey the observed signal. However, experts can lie and send $m \neq s$. A strategy of the expert is a function $\sigma(m|s,\hat{t})$ which gives the probability of sending a message m when the expert's information set is $\{\hat{t}, s\}$.

In the case of truthful revelation, $\sigma(m = s|s,\hat{t}) = 1$. We say that an expert uses a truthful strategy when $\sigma(m_i|s_i,\hat{t}) = 1,\ \forall i, \hat{t}$. We denote this truthful strategy σ^T.

Step 3 The agents observe m, make a conjecture $\tilde{\sigma}(m|s,\hat{t}),\ \forall s, \hat{t}$ on the expert's strategy, and update their belief on the state, $\nu_i(m,\tilde{\sigma},\tilde{t}) = p^a(\theta_i|m,\tilde{\sigma},\tilde{t})$.

Step 4 The updated prior of the agents determines the precision of the aggregate signal: if $\nu_1 \geq \chi^{-1}$, then $\tau(\nu_1)$ is equal to 1, and to 0 otherwise.

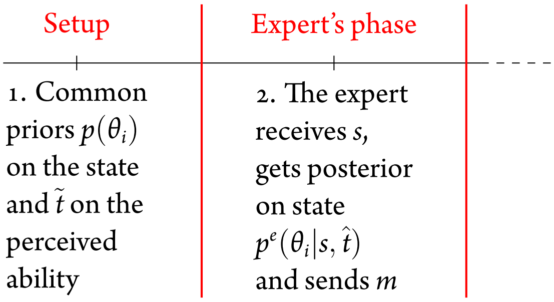

(a) Game with noisy observation of the state of the world

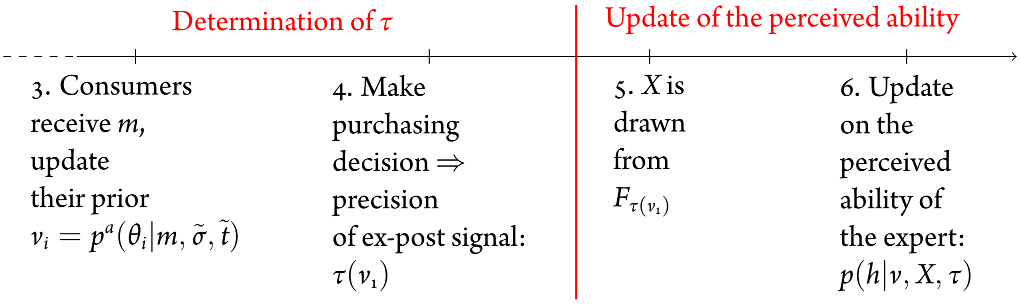

(b) Benchmark model

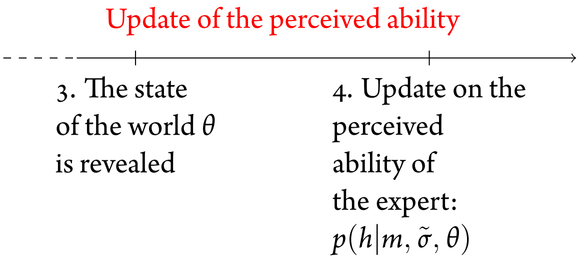

Figure 1.3.2: Timing of the game

Step 5 The realization of the aggregate signal X is drawn from its distribution $F_{\tau(v_1)}$.

Step 6 Agents compute the final posterior $p(\theta_i|v, X, \tau)$ and update their perception on the expert's ability $p(h|v, X, \tau)$.

1.3.3 Private posterior, updated prior, and final posterior

The expert and agents update their beliefs at each stage of the game using the Bayes rule.

Private posterior of the expert

$$p^e(\theta_i|s,\hat{t}) = \begin{cases} \frac{\hat{t}p(\theta_i)}{\hat{t}p(\theta_i)+(1-\hat{t})(1-p(\theta_i))} & \text{if } s = s_i \\ \\ \frac{(1-\hat{t})p(\theta_i)}{(1-\hat{t})p(\theta_i)+\hat{t}(1-p(\theta_i))} & \text{otherwise.} \end{cases} \tag{1.3.1}$$

Updated prior

Agents make a conjecture $\tilde{\sigma}$ on the strategy used by the expert. They can then form a belief over what signal the expert actually received $\tilde{s}$:

$$p^a(\tilde{s} = s|m, \tilde{\sigma}, \hat{t}) = \frac{\tilde{\sigma}(m|s,\hat{t})p(s,\hat{t})}{\tilde{\sigma}(m|s,\hat{t})p(s,\hat{t}) + \hat{\sigma}(m|s',\hat{t})p(s',\hat{t})}$$

The agents use their belief on $\tilde{s}$ to update their expectations on the state of the world, or update prior belief:

$$v_i = p^a(\theta_i|m,\tilde{\sigma},\tilde{t}) = \sum_{j\in\{0,1\}} p(\theta_i|s_j,\tilde{t}) \sum_{\hat{t}\in\hat{T}} p^a(\tilde{s}_j|m,\tilde{\sigma},\hat{t})p(\hat{t}|\tilde{t},m,\tilde{\sigma}) \tag{1.3.2}$$

where $p(\theta_i|s_i,\tilde{t})$ is defined as in Equation 1.3.1.

If the agents expect the expert to use a truthful strategy, the updated prior formula simplifies to:

$$p^a(\theta_i|m_j,\tilde{\sigma}^T,\tilde{t}) = p^a(\theta_i|\tilde{s}_j,\tilde{t})$$

Indeed, when the agents conjecture that the expert plays a truthful strategy σ^T, they necessarily believe the expert received a private signal s equal to her message m. In the following, when we consider that the agents expect a truthful strategy, we denote $\tilde{s}_i$ instead of m_i, $\tilde{\sigma}^T$.[4]

[4]In the remaining of the paper, we sometimes do not explicitly condition the probabilities on $\hat{t}$, $\tilde{t}$, or τ. In this case, it is

Final posterior

$$p^c(\theta_i|X,\tau,v_i) = \begin{cases} \frac{\tau(v_1)v_i}{\tau(v_1)v_i+(1-\tau(v_1))v_{-i}} & \text{if } X = X_i \\ \\ \frac{(1-\tau(v_1))v_i}{(1-\tau(v_1))v_i+\tau(v_1)v_{-i}} & \text{otherwise.} \end{cases} \tag{1.3.3}$$

1.4 Incentives to be Truthful and Structure of the Equilibrium

1.4.1 Definition of an equilibrium

We begin by introducing the definition of an equilibrium of the game.

Definition 1. *A perfect Bayesian equilibrium of the game is characterized by:*

1. *The expert chooses $\sigma^*(.|s,\hat{t})$ such that:*

$$\sigma^*(m|s,\hat{t}) = \arg\max_{\eta\in[0,1]} \eta E^e(p(h|m,\tilde{\sigma})|s,\hat{t}) + (1-\eta)E^e(p(h|m',\tilde{\sigma})|s,\hat{t})$$

2. *Agents correctly forecast the expert's strategy: $\tilde{\sigma}(.|s,\hat{t}) = \sigma^*(.|s,\hat{t})\ \forall s,\hat{t}$*

3. *Agents update their belief about the ability of the expert using Bayes' rule:*

$$p(h|m,\tilde{\sigma},X) = \sum_{\theta\in\{0,1\}} p(h|\theta,m,\tilde{\sigma})p(\theta|X,m,\tilde{\sigma})$$

In equilibrium, the agents' guess on the strategy of the expert is correct. In this game as in other cheap talk games, there is always a babbling equilibrium in which the expert sends uninformative messages. In the following sections, we are interested in informative equilibria in which the expert truthfully reveals her signal on some subset of prior beliefs. A necessary condition to have a informative equilibrium is that when the agents believe the expert plays a truthful strategy, the expert has indeed the incentive to reveal her private signal.

implicitly assumed that if a probability is conditional on $\tilde{s}$, it is also conditional on $\tilde{t}$. The same reasoning applies to s and $\hat{t}$. This allows us to substantially lighten the notations.

1.4.2 Incentive Compatible Sets and Truthful Revelation Sets

Definition of the Incentive Compatible Sets

To derive the features of the most informative equilibrium, we start by assuming that the agents believe experts of all abilities plays truthfully and we identify the priors for which it is optimal for the expert to play according to the agents' belief. The expert has the incentive to truthfully reveal her private signal if she expects the audience to perceive her as more able when she reveals her signal than when she lies and agents do not anticipate it:

$$p(h|\tilde{s}_i, s_i, \tilde{t}, \hat{t}) \geq p(h|\tilde{s}_{-i}, s_i, \tilde{t}, \hat{t}) \tag{1.4.1}$$

Solving Equation 1.4.1 for $i = 0, 1$ allows us to recover the range of prior beliefs over which the expert has an incentive to play truthfully. This gives the following definition:

Definition 2. *The incentive compatible set of an expert with (perceived or true) ability $\hat{t}$ and perceived ability $\tilde{t}$ is denoted $IC_{\hat{t},\tilde{t}}$ and defined as:*

$$IC_{\hat{t},\tilde{t}} = \left\{\mu_1 \in [0, 1] \mid p(h|\tilde{s}_i, s_i, \tilde{t}, \hat{t}) \geq p(h|\tilde{s}_{-i}, s_i, \tilde{t}, \hat{t}),\ \forall i\right\}$$

We can rewrite the condition determining if a prior belongs to the incentive compatible set of the expert.

We denote p^+ and p^- the positive and negative updates on the perceived ability when the agents observe the true state of the world at the end of the game, namely:

- $\forall i,\ p^+ = p(h|\hat{s}_i, \theta_i) = \frac{hp(h)}{hp(h)+lp(l)}$
- $\forall i,\ p^- = p(h|\hat{s}_{-i}, \theta_i) = \frac{(1-h)p(h)}{(1-h)p(h)+(1-l)p(l)}$

Equation 1.4.1 rewrites:

$$\begin{aligned}
&p(h|\tilde{s}_i, s_i, \tilde{t}, \hat{t}) \geq p(h|\tilde{s}_{-i}, s_i, \tilde{t}, \hat{t}) \\
\Leftrightarrow &p^+ p^e(\theta_i^a|\tilde{s}_i, s_i) + p^- p^e(\theta_{-i}^c|\tilde{s}_i, s_i) \geq p^+ p^e(\theta_i^a|\tilde{s}_{-i}, s_i) + p^- p^e(\theta_{-i}^a|\tilde{s}_{-i}, s_i) \\
\Leftrightarrow &(p^+ - p^-)p^e(\theta_i^a|\tilde{s}_i, s_i) + p^- \geq (p^+ - p^-)p^e(\theta_{-i}^a|\tilde{s}_{-i}, s_i) + p^- \\
\Leftrightarrow &p^e(\theta_i^a|\tilde{s}_i, s_i) \geq p^e(\theta_{-i}^a|\tilde{s}_{-i}, s_i)
\end{aligned} \tag{1.4.2}$$

where the second equivalence stems from $p^+ > p^-$.[5] Inequality 1.4.2 says that the expert has the incentive to be truthtelling if and only if she is more credible when telling the truth than when lying.

[5]For non trivial values of parameters.

Formally, after receiving s_i, the agents' final posterior on θ_i *as expected by the expert* and given m_i must be higher than the agents' final posterior on θ_{-i} given m_{-i}.

Definition of the Truthful Revelation Set

The incentive compatible sets are computed assuming agents believe experts of all abilities are truthtelling. They are useful to study the equilibrium. In particular, if a given prior belief belongs to the incentive compatible sets of all experts, then truthtelling is an equilibrium when an expert faces this prior belief.

However, truthtelling might be an equilibrium for experts of certain abilities on prior beliefs falling outside their incentive compatible sets. For instance, a high ability expert who is aware about her ability might want to signal her ability to the agents by truthfully reporting contradictory signals, even for extreme prior beliefs. Doing so, she separates from less able experts who cannot mimic her strategy without hurting seriously how the audience perceives them.

We define the truthful revelation set as the set of prior beliefs for which truthful revelation of the private signal is an equilibrium for experts of all abilities. It is the intersection of the incentive compatible sets of experts of all abilities.

Definition 3. *The truthful revelation set is the subset of prior beliefs for which truthful revelation of the private signal is an equilibrium for experts of all abilities. It is given by:*

$$TRS_{\tilde{t}} = \cap_{\hat{t}} IC_{\hat{t},\tilde{t}}$$

The truthful revelation set is an indicator on how informative the equilibrium actually is. When the expert does not know her own ability, i.e. $\hat{t} = \tilde{t}$, there is only one incentive compatible set and the truthful revelation set coincides with it.

Computation of the Incentive Compatible Sets

With the *All or Nothing* τ function for the precision of the aggregate signal, agents observe the true state if their updated prior rises above the threshold and an uninformative aggregate signal otherwise. In that case, their updated prior constitutes the only piece of information they have on the state of the world and they use it to update their perception on the expert's ability.

The expert truthfully reports her private signal if she thinks agents will see her as able, i.e. as an expert who got her signal right. If an expert thinks the agents are unlikely to believe that the state of the world corresponds to her signal, she does not reveal it.

With the *All or Nothing* τ function for the precision of the aggregate signal, the expectation of the expert on the ultimate belief of the agents depends on the updated prior ν_1 and the threshold χ^{-1}. If the updated prior falls below the threshold, agents only observe an uninformative aggregate signal and their belief is the updated prior ν. If the updated prior rises above the threshold, agents will observe the true state, which the expert expects to be distributed according to $p^e(\theta_i|s,\hat{t})$.

Formally:

$$p^e(\theta_i^a|\tilde{s},s,\tilde{t},\hat{t}) = \begin{cases} \nu_i(\tilde{s},\tilde{t}) & \text{if } \nu_i(\tilde{s},\tilde{t}) < \chi^{-1} \\ p^e(\theta_i|s,\hat{t}) & \text{if } \nu_i(\tilde{s},\tilde{t}) \geq \chi^{-1} \end{cases}$$

The expert faces three different possibilities.

(A) Whichever message she sends, the initial prior is so low that the updated prior of the agents falls below the threshold. As the updated prior is always higher when agents believe the expert received a high signal, case (A) corresponds to $\nu_1(\tilde{s}_1,\tilde{t}) < \chi^{-1}$.

(B) Whichever message she sends, the initial prior is so high that the updated prior of the agents is greater than the threshold. As the updated prior is always lower when agents believe the expert received a low signal, case (B) corresponds to $\nu_1(\tilde{s}_0,\tilde{t}) > \chi^{-1}$.

(C) The initial prior is intermediate, and the updated prior of the agents either falls below or above the threshold, depending on whether the agents believe the expert has received a low or a high signal. Case (C) corresponds to $\nu_1(\tilde{s}_0,\tilde{t}) \leq \chi^{-1} \leq \nu_1(\tilde{s}_1,\tilde{t})$.

The areas of prior beliefs that determine whether or not the expert can manipulate the observability of the state of the world depend on her perceived ability. Figure 1.4.1 depicts how the areas change with the perceived ability for $\chi^{-1} = \frac{1}{4}$. As the (C) area expands with the perceived ability, the expert with a higher perceived ability has a greater power to manipulate.

In case (A), agents only use their updated prior to update their perception about the expert's ability. As a result, the expert always has incentives to confirm the initial prior μ_1. The incentive compatible set boils down to $\{\frac{1}{2}\}$. Formally, the expert has the incentive to truthfully reveal her signal if and only if:

$$\nu_i(\tilde{s}_i,\tilde{t}|s_i) \geq \nu_{-i}(\tilde{s}_{-i},\tilde{t}|s_i) \quad \forall i \Leftrightarrow \mu_i \geq \frac{1}{2} \quad \forall i$$

In case (B), the agents observe the true state of the world at the end of the game whatever the action of the expert. We are back to Ottaviani and Sørensen [41] or Ottaviani and Sørensen [42]. The expert wants to truthfully reveal her signal whenever she is confident it is a valuable piece of information, i.e. whenever her signal does not contradict her posterior on the state of the world. Formally, the expert has the incentive to truthfully reveal her signal if and only if:

$$p^e(\theta_i|s_i,\hat{t}) \geq p^e(\theta_{-i}|s_i,\hat{t}) \quad \forall i \Leftrightarrow 1-\hat{t} \leq \mu_1 \leq \hat{t}$$

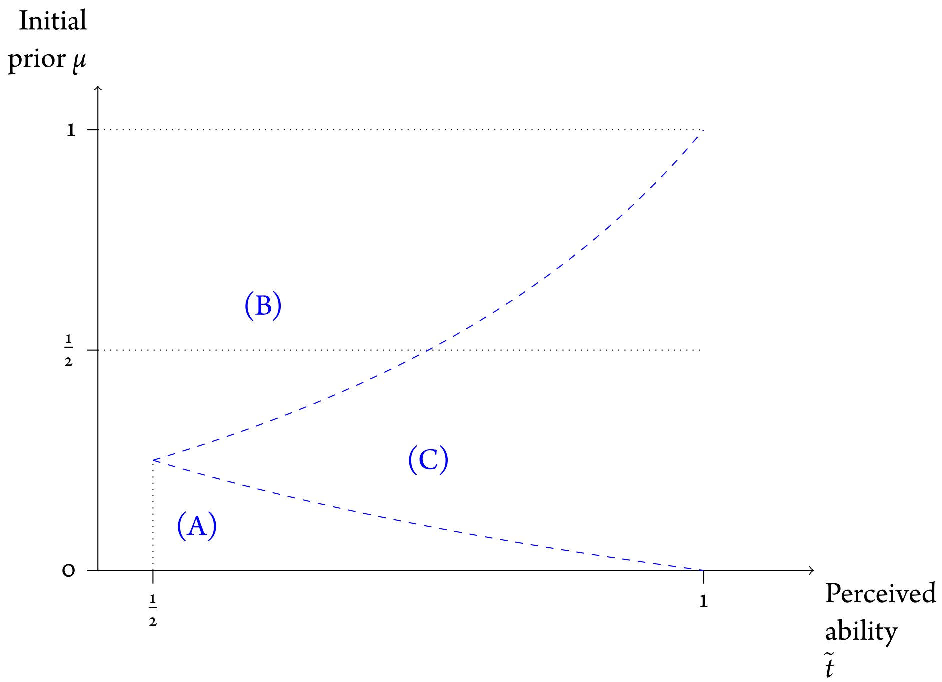

Figure 1.4.1: The manipulation of the observability of the state of the world is possible in area (B) — Areas drawn for $\chi^{-1} = \frac{1}{4}$

When the initial prior is intermediate however, i.e. in case (C), the expert can influence the precision of the aggregate signal through her message. An initial prior on the high state of the world μ_1 belongs to the incentive compatible set of the expert if the expert wants to reveal truthfully both a low signal and a high signal, which either leads to an uninformative aggregate signal or the perfect revelation of the state of the world.

The expert has the incentive to truthfully reveal a low signal if and only if:

$$v_0(\tilde{s}_0, \tilde{t}) \geq p^e(\theta_1 | s_0, \hat{t}) \Leftrightarrow \begin{cases} \mu_1 \leq \tilde{t} & \text{when the expert does not know her own ability,} \\ \mu_1 \leq \frac{1}{1+\sqrt{\frac{1-t}{t}\frac{1-\tilde{t}}{\tilde{t}}}} & \text{otherwise.} \end{cases}$$

The expert has the incentive to truthfully reveal a high signal if and only if:

$$p^e(\theta_1 | s_1, \hat{t}) \geq v_0(\tilde{s}_0, \tilde{t}) \Leftrightarrow \begin{cases} \mu_1 \geq \frac{1}{2} & \text{when the expert does not know her own ability,} \\ \mu_1 \geq \frac{1}{1+\sqrt{\frac{t}{1-t}\frac{1-\tilde{t}}{\tilde{t}}}} & \text{otherwise.} \end{cases}$$

1.4.3 The Most Informative Equilibrium with Unknown Ability

In the case of unknown ability, the expert and the agents share a common belief about the expert's ability and $\hat{t} = \tilde{t}$. In that case, the truthful revelation set coincides with the incentive compatible set.

Proposition 1. *The most informative equilibrium displays truthful revelation within the bounds of the incentive compatible set $IC_{\tilde{t}}$ and babbling outside.*

Proof. By definition, outside the incentive compatibility set, the expert is strictly better off by lying.

This excludes any possibility of a mixed equilibrium. The expert cannot send any credible message outside $IC_{\tilde{t}}$. Obviously, for all $\mu_1 \in IC_{\tilde{t}}$, truthtelling is an equilibrium. □

The shape of the equilibrium in the case of unknown ability is simple. It does not depend on the assumption we made about τ. The expert cannot credibly transmit any valuable information outside the truthful revelation set.

1.4.4 The Most Informative Equilibrium with Privately Known Ability

When the expert knows her own ability, i.e. $\hat{t} = t$, the truthful revelation set is the intersection of the incentive compatible sets of the low and the high ability experts.

Increasing Incentive Compatible Sets

When experts know their abilities, the incentive compatible sets of experts with low abilities are included in the incentive compatible sets of experts with high abilities. In other words, when a low ability expert is confident enough to tell the truth, a high ability expert is also confident enough. This result holds for a given, fixed, perceived ability. Obviously, able experts should be able to gain the audience's trust on their ability over time.

Proposition 2. *$IC_{l,\tilde{t}} \subseteq IC_{h,\tilde{t}}$ for $\frac{1}{2} < l < h \leq 1$.*

We detail the proof of this result in Appendix A.2. The proof does not rely on the assumption we made on τ and the result extends to the general case.

A consequence of the inclusion of incentive compatible sets is that the truthful revelation set is equal to the incentive compatible set of the expert with the lowest ability, here l.

Structure of the Equilibrium

When the experts know their own ability, the equilibrium separates high and low ability experts on extreme prior beliefs.

Proposition 3. *With known abilities, the most informative equilibrium of the game is characterized by the following strategy:*

- *The high ability expert is truthtelling over the entire prior space.*
- *The low ability expert is truthtelling over her incentive compatible set $IC_{l,\tau} = \left[\underline{\mu}_1, \overline{\mu}_1\right]$. For $\mu_1 < \underline{\mu}_1$, the expert sends m_0 after observing s_0 and sends m_1 with a positive probability after observing s_1. For $\mu_1 > \overline{\mu}_1$, the expert sends m_1 after observing s_1 and sends m_0 with a positive probability after observing s_0.*

We give a detailed proof in Appendix A.3. The proof relies on the result on the inclusion of incentive compatible sets and is designed along the lines of Ottaviani and Sørensen [41]. The proof of Proposition 3 also extends to a general τ function.

1.5 The Barrier to Entry Effect

If the prior belief is to low or too high, i.e. if the prior belief falls in the area (A) or the area (B), the incentive compatible set of the expert either reduces to $\left\{\frac{1}{2}\right\}$ or equals the one in models with perfect revelation of the state of the world. In both cases, the incentive compatible set is symmetric around the average prior belief $\frac{1}{2}$. Even though experts can misreport their signals in those cases, the misreporting is as likely for states of the world with high priors than states of the world with low priors.

If the prior belief falls in area (C) however, i.e. when the expert can manipulate the observability of the state of the world, the incentive compatible set of the expert shifts towards high prior beliefs, as shown by the bounds derived in section 1.4.2. We call *barrier to entry effect* this shift of the incentive compatible set towards high prior beliefs.

In area (C), the range of priors defining the incentive compatible set of the expert not knowing her ability is:

$$\frac{1}{2} \leq \mu_1 \leq \tilde{t}$$

When the expert knows her ability, the range of priors defining the incentive compatible set of the expert is:

$$\frac{1}{1+\sqrt{\frac{t}{1-t}\frac{1-\tilde{t}}{\tilde{t}}}} \leq \mu_1 \leq \frac{1}{1+\sqrt{\frac{1-t}{t}\frac{1-\tilde{t}}{\tilde{t}}}}$$

Figure 1.5.1 represents the asymmetric incentive compatible sets of the expert, should the precision of the aggregate signal pass from $\frac{1}{2}$ to 1 when the expert changes her message from 0 to 1.

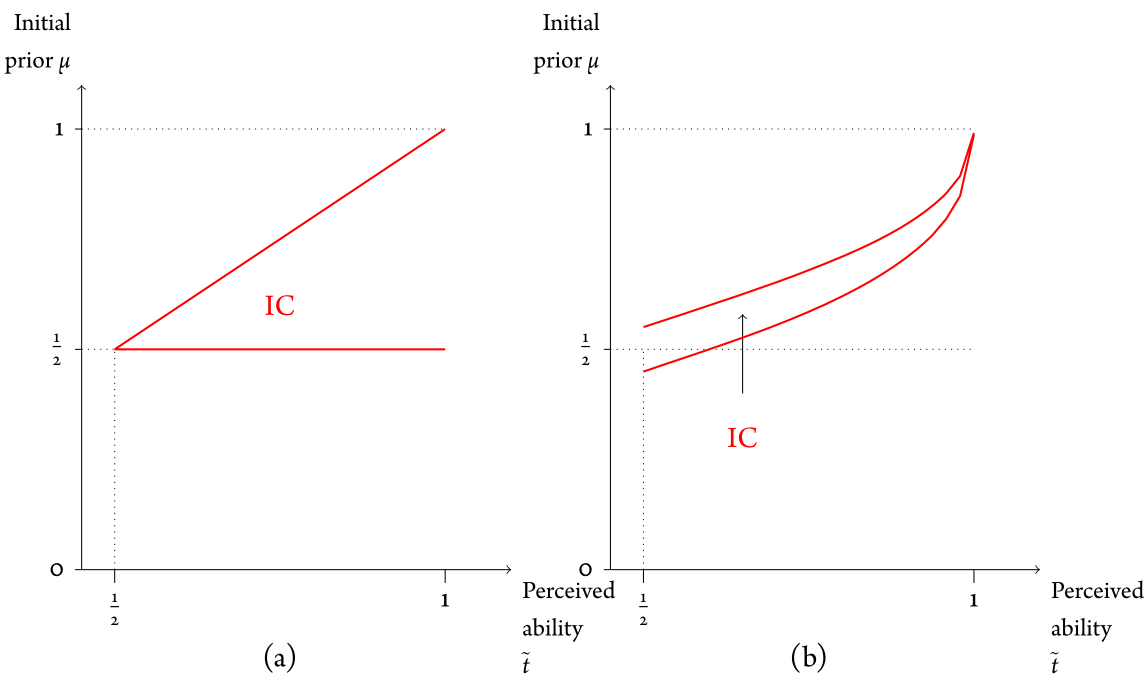

Figure 1.5.1: Barrier to entry effect when the expert's ability is unknown (a) or known by the expert (b) — The incentive compatible set when the expert knows her ability is drawn for $t = 0.6$.

When the expert does not know her own ability and can manipulate the observability of the state of the world through her message, no prior belief on the high state of the world below $\frac{1}{2}$ belongs to the incentive compatibility set of the expert.

When the expert knows her ability and can manipulate the observability of the state of the world through her message, the incentive compatible set of the expert is strictly above priors equal to $\frac{1}{2}$ as soon as the perceived ability is greater than the true ability, $\tilde{t} > t$. Experts whose perceived ability is greater than their true ability are not enticed to truthfully reveal their private information even when the state of the world is ex-ante as likely to be low than to be high.

Hence, as long as the expert can manipulate the observability of the state of the world, the misreporting of private information is more likely for states of the world with low prior beliefs than for states of the world with high prior beliefs. Moreover, the barrier to entry effect is more pronounced when the expert knows her ability.[6]

1.6 The Guru Effect

Allowing the precision of the aggregate signal to depend on the updated prior belief of the audience yields a second effect that is absent from reputational cheap talk models where the precision of the aggregate signal is exogenous : the *guru effect*. The guru effect arises when for some low prior belief on the high state of the world, an expert with a low perceived ability reveals truthfully a high private signal, but an expert with a higher perceived ability wants to use her influence and to manipulate the observability of the state of the world. She has the incentive to send the low message to prevent the audience to update her ability. In equilibrium, the expert with the low perceived ability is credible, but not the expert with the high perceived ability. We give a formal definition for the guru effect.

Definition 4. *The guru effect arises when:*

$$\exists \mu_1 \in [0,1] \text{ and } [\tilde{t}_1, \tilde{t}_2],\ s.t.\ \mu_1 \in IC_{\hat{t},\tilde{t}} \quad \forall \tilde{t} \in [\tilde{t}_1, \tilde{t}_2] \text{ but } \mu_1 \notin IC_{\hat{t},\tilde{t}'} \quad \forall \tilde{t}' > \tilde{t}_2$$

1.6.1 Guru Effect with Unknown Ability

For some prior beliefs, the expert can manipulate the precision of the ex-post aggregate signal. (Assuming that the audience trusts her.) The area (C) depicted in Figure 1.4.1 represents such priors. For other prior beliefs, the expert has the incentive to manipulate the precision of the ex-post aggregate signal to maintain

[6] Only for the low ability expert as the cut-off at which the incentive compatible set goes above one half increases in t. Moreover, in equilibrium, the high ability expert truthfully reveals her signal on the whole range of prior beliefs.

the audience's belief on her ability, creating the barrier to entry effect and the asymmetric incentive compatibility set depicted in Figure 1.5.1 (a).

For some values of the threshold above which the precision of the aggregate signal is 1, an expert with a low perceived ability knows that even if she sends a negative recommendation, she will not be persuasive enough to prevent the audience from observing the true state of the world. But if her perceived ability increases, she can persuade the audience not to observe the true state and the incentive to send a negative recommendation dominates. This incompatibility between the will to manipulate the precision of the aggregate signal and the incapability to do so when the perceived ability is low creates the guru effect.

Paradoxically, the gains in credibility that should come from a higher perceived ability give perverse incentives to the expert who loses all credibility in equilibrium as agents anticipate her strategy. With the guru effect, experts with lower perceived abilities are more credible for some range of prior beliefs.

The guru effect only arises when it is difficult for the low ability expert to convince the audience not to observe the true state of the world. (Or when the mechanism through which individual signals aggregate to the ex-post aggregate signal is efficient.)

Proposition 4. *The guru effect only arises for $\mu_1 < \frac{1}{2}$ and when $\chi^{-1} < \frac{1}{2}$.*

Proof. To prove the first part of the Proposition, suppose we have a prior $\mu_1 \geq \frac{1}{2}$ and a perceived ability $\tilde{t}_1$ such that an expert with that perceived ability is truthful. Then we necessary have $v_1(\tilde{s}_1, \tilde{t}_1) \geq \chi^{-1}$, otherwise the expert would not be truthful, i.e. we are in areas (B) or (C). As $\mu_1 \geq \frac{1}{2}$ and the expert is truthful, looking at the incentive compatible conditions for (B) and (C), we necessarily get $\tilde{t}_1 \geq \mu_1 \geq \frac{1}{2}$, which implies $\tilde{t}_2 \geq \mu_1 \geq \frac{1}{2}$ and necessary truthtelling for the expert $\tilde{t}_2$. Hence, the guru effect can only arise for $\mu_1 < \frac{1}{2}$.

Now, suppose we have a prior $\mu_1 < \frac{1}{2}$ and a perceived ability $\tilde{t}_1$ such that an expert with that perceived ability is truthful. Then we are necessarily in area (B) and $1 - \tilde{t}_1 \leq \mu_1$ and $\mu_1 \geq \frac{\tilde{t}_1}{1+(\chi-2)(1-\tilde{t}_1)} \equiv f(\chi, \tilde{t}_1)$. (Conditions for the incentive compatible set and for being in area (B).) As $f(\chi = 2, \tilde{t}_1) = \tilde{t}_1$ and $\frac{\partial f(\chi,\tilde{t}_1)}{\partial \chi} < 0, f(\chi, \tilde{t}_1) > \tilde{t}_1$ for $\chi < 2$. But then $\mu_1 < \frac{1}{2}$ is incompatible with $\mu_1 \geq f(\chi, \tilde{t}_1)$ for $\chi^{-1} \geq \frac{1}{2}$.

If $(\mu_1, \tilde{t}_2)$ still belongs to (B), the expert with perceived ability $\tilde{t}_2$ is necessarily truthtelling. Hence, to have a guru effect $\tilde{t}_2$ must be high enough to switch in areas (C) or (A), which means $v_1(\tilde{s}_0, \tilde{t}_2) \leq \chi^{-1} \Leftrightarrow \mu_1 \leq \frac{\tilde{t}_2}{1+(\chi-2)(1-\tilde{t}_2)}$. Proving that such a $\tilde{t}_2$ exists for $\chi^{-1} < \frac{1}{2}$ is straightforward. □

Hence, the guru effect is more likely when only the reports of a few agents in the audience are sufficient to create a fairly precise aggregate signal. Figure 1.6.1 represents the truthful revelation set when $\chi^{-1} = \frac{1}{4}$, which displays a significant guru effect. (When the expert does not know her own ability, the incentive compatible set is the truthful revelation set.)

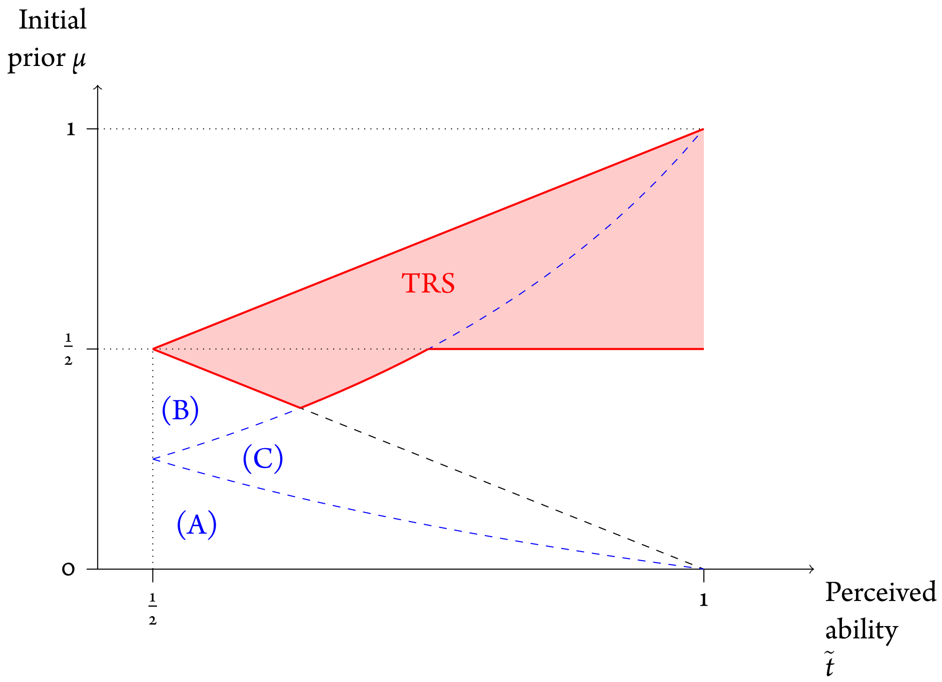

Figure 1.6.1: The guru effect appears in the truthful revelation set when $\chi^{-1} < \frac{1}{2}$. (Here $\chi^{-1} = \frac{1}{4}$.)

1.6.2 Guru Effect with Known Ability

When the expert knows her ability, the guru effect appears already in the shape of the incentive compatible set in area (C), i.e. in the area where the expert can manipulate the precision of the aggregate signal. In that case, the range of prior beliefs compatible with truthful revelation of the signal shifts upward as the perceived ability increases. Figure 1.5.1 presents an illustration of the guru effect.

Contrary to the situation with unknown ability, the guru effect is directly visible in the incentives of the expert. As the barrier to entry effect, the guru effect is more pronounced in the case where the expert knows her ability. But in equilibrium, it is counterbalanced by the even stronger incentives for the able expert to separate from the low ability one, and for the low ability one to mimic the able expert.

However, when we compute the truthful revelation set, which we present in Figure 1.6.2, truthtelling is never incentive compatible for the low ability expert, for some values of the perceived ability.

To compute the truthful revelation set, we add the incentive compatible conditions to the conditions

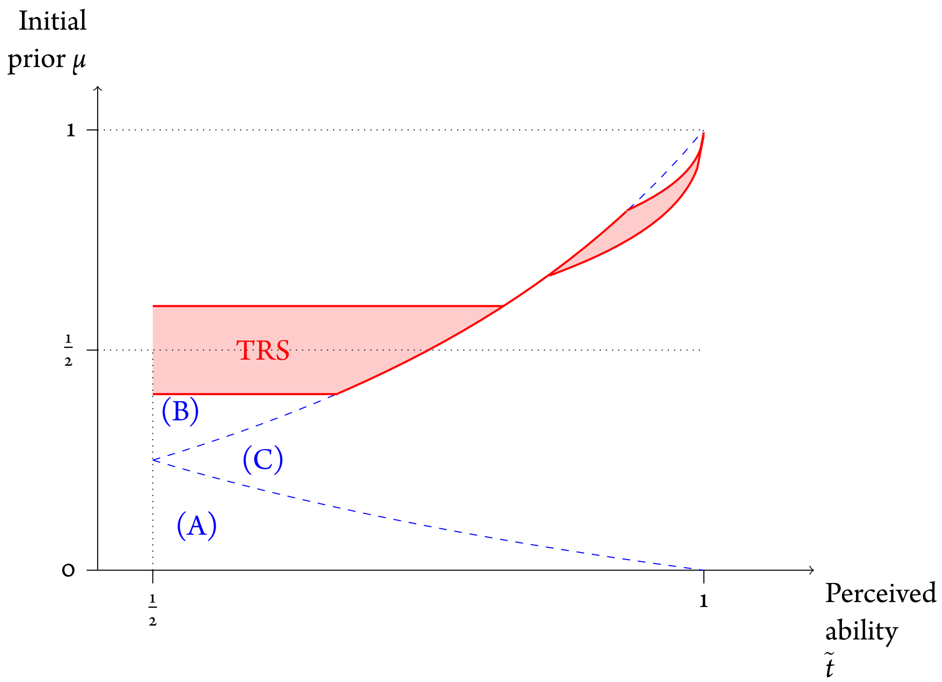

Figure 1.6.2: When the expert knows her ability, some values of her perceived ability annihilate incentives to be truthtelling. (Here $\chi^{-1} = \frac{1}{4}$ and $l = 0.6$.).

for being able to manipulate the precision of the aggregate signal, which gives:

$$IC_{t,\tilde{t}} = \begin{cases} 1-t \leq \mu_1 \leq t & \text{if } \nu_1(\tilde{s}_0, \tilde{t}) > \chi^{-1} \\ \frac{1}{1+\sqrt{\frac{t}{1-t}\frac{1-\tilde{t}}{\tilde{t}}}} \leq \mu_1 \leq \frac{1}{1+\sqrt{\frac{1-t}{t}\frac{1-\tilde{t}}{\tilde{t}}}} & \text{if } \nu_1(\tilde{s}_0, \tilde{t}) \leq \chi^{-1} \leq \nu_1(\tilde{s}_1, \tilde{t}) \\ \mu_1 = \frac{1}{2} & \text{if } \nu_1(\tilde{s}_1, \tilde{t}) \leq \chi^{-1} \end{cases}$$

Such failures to elicit incentives to be truthtelling on some range of priors and perceived ability only arise if $\chi^{-1} < \frac{1}{2}$. Also, we do not find empty truthful revelation sets as soon as the function τ determining the precision of the aggregate signal is smooth.

1.7 Extension to the General Case

We give a brief description of the previous results that extend to a general τ function. When the τ function is increasing (and smooth), we can express the agents' posterior belief as expected by the expert $p^e(\theta^a|\tilde{s}, s, \tilde{t}, \hat{t})$ as a weighted average between the expert's private posterior and the agents' updated prior. This is a more general form for the characterization of the incentive compatible set.

1.7.1 Expert not Knowing her Ability

We extend the results on the barrier to entry effect when the expert does not know her own ability. We also get a preliminary result on the guru effect in this case.

We also show that the extent of information revealed in equilibrium increases with the precision of the aggregate signal. For two τ functions, if one is above the other, the truthful revelation set associated to this function is larger. It means that fast mechanisms of information transmission in the pool of agents yield a higher degree of information transmission by the expert. Applied to our simplified model, the truthful revelation set shrinks when the value of the threshold increases. It also means that the greater degree of information transmission in the simplified model comes from the guru effect.

Proposition 5. *The truthful revelation set is increasing with the precision of the aggregate signal: Take two functions* $\tau_1, \tau_2 : [0,1] \to [1/2, 1]$*, such that* $\forall x \in [0,1]\,,\ \tau_1(x) < \tau_2(x)$*, then* $IC_{\tilde{t},\tau_1} \subseteq IC_{\tilde{t},\tau_2}$.

We detail all the proofs of the generalized results in the Appendix A.4.

1.7.2 Expert Knowing her Ability

When the expert knows her ability, the bounds $\underline{\mu}_1, \overline{\mu}_1$ of the truthful revelation set (i.e. the incentive compatible set of the low ability expert) now depend on the shape of the function $\tau(.)$ and determining their behaviour according to the initial perceived ability becomes a difficult exercise. However, numerical examples seem to indicate that our results on the barrier to entry and guru effect extend to this framework as well.

1.7.3 Competition between experts

Another extension of our model is the competition between experts. When the precision of the aggregate signal is endogenous, competition creates more scope to learn the true state of the world and hence limits the value of manipulating the signal in order to change the precision.

Competition has an positive impact on the amount of information revealed as it softens the incentives to manipulate the precision of the aggregate signal. This contrasts with the result of Ottaviani and Sørensen [42] who show that when the true state of the world is ultimately revealed, competition between experts does not matter.

1.8 Conclusion

Introducing endogeneity of the precision of the aggregate signal distorts the incentives of the expert to reveal truthfully her private information. Compared to models in which the audience observes the state of the world perfectly, the expert has incentives to strategically manipulate the information she sends to her audience: even if she deeply believes that a high state of the world is more likely, she might want to send a negative recommendation. We recover two additional effects: the barrier to entry effect and the guru effect.

The endogeneity of the precision of the aggregate signal also yields a positive impact of competition on information transmission, a feature missing form reputational cheap-talk games with perfect observability of the state of the world.

In our model, when the threshold above which the state of the world becomes perfectly observed increases, information transmission by the expert decreases. Hence, experts have less incentives to truthfully reveal their opinions when the audience faces a higher cost of observing the true state of the world, i.e. when the audience would benefit the most from the honest advice of an expert.

Overall our work could help to assess the reliability of experts reports given their past experience and influence. The effects we recover should be taken into account in empirical analyses of expert reviews, if the experts have some power on the observability of the state of the world.

On that point, Socrates, I have heard that one who is to be an orator does not need to know what is really just, but what would seem just to the multitude who are to pass judgment, and not what is really good or noble, but what will seem to be so; for they say that persuasion comes from what seems to be true, not from the truth.

Plato, *Phaedrus*, 260a

2

Structural Estimation of Expert Strategic Bias: the Case of Movie Reviewers

co-authored with Fanny Camara

Abstract

We develop the first structural estimation of reputational cheap-talk games using data on movie reviews released in the US between 2004 and 2013. We identify and estimate movies' priors, as well as movie reviewers' abilities and strategic biases. We find that reviewers adopt reporting strategies that are consistent with the predictions of the literature on reputational cheap-talk. The average conservatism bias for low prior movies lies between 8 and 11%, depending on the specifications of the model. The average conservatism bias for high prior movies ranges from 13 to 15%. Moreover, we find a significant, albeit small, effect of the reputation of the reviewers on their strategies, indicating that incentives to manipulate demand in order to prevent reputation updating are present in this industry. Our estimation takes into account and quantifies potential conflicts of interest that might arise when the movie reviewer belongs to the same media outlet as the film under review. Out-of-sample predictions confirm that movie reviewers do have reputational concerns.

2.1 Introduction

We use a reputational cheap talk framework to model the reviewing behavior of movie critics. Cheap-talk games describe situations in which a privately informed expert sends a prediction about the state of the world to some receivers. The expert sends her recommendation at no cost and cannot base her report on any verifiable (hard) evidence. In this paper, we focus on the subclass of reputational cheap-talk games in which the expert seeks to maximize her reputation for being well informed. The expert's reputation corresponds to the receivers' belief about the precision of her private information. We develop a new structural approach to quantify the strategic biases of movie reviewers which arise from their reputational concerns. Our approach allows us to identify and estimate individual reviewers' abilities and biases. Our model has a stronger predictive power than alternative models without reputational concerns, which gives us empirical evidence that career concerns do shape the behavior of movie reviewers.

Movie reviewers watch films before their official release and get private signals on their quality whose precision depends on their ability. The career of these reviewers is built on their reputation for accuracy. Any expert assessing the quality of experience goods has similar reputational concerns. Financial analysts sell their advice on investment opportunities about which they are supposed to have some private information. A better reputation for good advice allows them to charge higher prices. Academic referees in peer-reviewed scientific journals give their opinion on whether or not a paper is publishable based on their personal assessment of its scientific quality. Signalling their good judgement to the editor can help their academic career.

In all those settings, the theoretical literature shows that experts may strategically misreport their private information. In their seminal paper on reputational cheap-talk, Ottaviani and Sørensen [42] show that experts may disregard noisy signals and conform to the prevalent opinion in order to pass for good predictors of the state of the world. Hence, when there is a strong prior belief on the state of the world being high and the expert receives a private signal supporting the opposite, she has the incentive to lie and to pretend that she received a high signal. This tendency for the expert to stick to extreme priors creates a conservatism bias. In the most informative equilibrium, despite this incentive, the high ability expert truthfully reveals her signal while the unskilled one conforms to the common prior with some positive probability. In this binary state setting, when the public belief about the state is more balanced, both types of experts are inclined to truthfully reveal their signals.

Camara and Dupuis [11] and Mariano [34] study a similar game in which the state of the world is imprecisely revealed. These papers endogenize the precision with which receivers can observe the state after its realization. In particular, the expert can make it harder for the receivers to observe the true state by convincing them that it is low. For instance, movie reviewers can discourage consumers to see a movie by writing harsh reviews. It is even more true for obscure movies. In this case, the scarce audience of the movie prevents the market from learning about its quality. Advice not to invest lead to no investment and thus to the failure of the project, regardless of its quality, which then becomes unobservable. When unsure about the accuracy of her high signal, the expert has the incentive to send a low report in order to garble her reputation update. This extra incentive for experts to over-report the low state generates what we call a manipulation bias. This manipulation bias is stronger for influential experts, i.e. experts who enjoy a high reputation, as their recommendations are more likely to change drastically receivers' beliefs. Contrary to the case of perfect ex-post revelation of the state of the world, both the reputation of the expert and the polarity of the prior, i.e. whether it is low or high, affect the experts' strategies.

In this paper, we develop the first structural estimation of cheap-talk games and apply it to movie reviews. We estimate the unobserved ability and quantify the extent of misreporting for each expert in our sample. Moreover, our estimation takes into account and quantifies potential conflicts of interest that might arise when the movie reviewer belongs to the same media outlet as the film under review.

The challenge in estimating biases in the reputational cheap-talk game played by movie reviewers stems from the unobservability of the movies' true quality, the experts' abilities, and of their private signals. Our estimation strategy has the double advantage not to rely on an estimation or approximation of quality to which we would compare the reviews, and to allow us to recover the ability and prior-dependent biases of each expert in our sample. The estimation directly exploits the structure of our theoretical model.

Following Camara and Dupuis [11], we extend the model to a continuum of expert's abilities and provide a simple characterization of the incentives to be truthtelling. This allows us to characterize a unique set of priors over which every expert, regardless of her ability, is truthtelling in the most informative equilibrium. This truthtelling set is at the core of our estimation strategy. In our empirical analysis, we control for horizontal differentiation by using data on MPAA ratings which are parental advisory guidelines that are highly correlated with genres. The quality is the vertical element of differentiation between movies: within each genre, a high quality movie is more likely to please a consumer than a low quality movie. We identify movies priors as well as experts' biases, and abilities. For given biases and abilities, higher prior movies tend to get more unanimously good reviews. For a given prior, we retrieve both abilities and biases using the partition of the prior space given by the equilibrium conditions of our theoretical model. This partition consists in priors falling within or outside a truthful revelation set. This set is the subset of the prior space on which experts of any ability are truthtelling in equilibrium. Its position in the prior space also depends on the reputation of the expert. Inside the truthful revelation set, absent of bias, the distribution of observed recommendations is determined by priors and experts' abilities. We can therefore recover abilities for given priors. For given priors and abilities, we are able to simulate the distribution of reviews that critics would give if they were truthful and compared this distribution with the one we observe for movies whose priors fall outside the truthful revelation set. The distance between those two distributions gives us the expert's bias. We observe variation over time in the reputation of each reviewer which shifts their truthful revelation set and allow us to identify its effect on each reviewer's bias. Our estimation strategy directly proceeds from the identification since we use the same partition structure over the set of priors in a maximum likelihood approach.

Our estimation strategy relies on weak assumptions: we assume that, for a given prior, experts of all abilities report truthfully their private signals if such an equilibrium exists. An alternative for the movie reviewers would be to play a babbling equilibrium in which their recommendations are completely uninformative. Since we observe that movie reviewers send reviews correlated to the true quality, we think this is highly unlikely.[1] Moreover, as shown by Basuroy et al. [4], Reinstein and Snyder [45], or Basuroy and Ravid [3], consumers tend to take into account movie reviews when they make their purchasing decision. This behavior is inconsistent with a babbling equilibrium. Our estimation of the bias does not rely on any additional assumptions concerning which equilibrium is played outside the truthtelling set, which is fortunate considering that cheap-talk games usually feature multiple equilibria.

We collect a nearly exhaustive data set on movies released in the US between 1990 and 2013 that contains the main information on directors, production companies, budgets, exact release dates, and genres. For each of those movies we gather the reviews published on *rottentomatoes.com*, the reference website displaying critics from the most influential movie reviewers. The reviews posted on this website have a binary structure which perfectly matches our theoretical model. (The movie is deemed *fresh* or *rotten*.) Along with the name of the reviewer and her positive or negative recommendation, we collect the medium in which the review was published and the date of the publication. We use *Google trend* to retrieve monthly data on the number of Google searches which measure reviewers' reputation.

We find movie reviewers' abilities ranging from 62% to 90%, meaning that the least able movie reviewer in our sample receives the correct private signal 6 times out of 10 whereas the most able one is correct 9 times out of 10. The average conservatism bias for low prior movies, i.e. the average probability that the experts in our sample transform a positive signal into a bad review when the prior is low is between 8 and 11%, depending on the specifications of the model. The average conservatism bias for high prior movies ranges from 13 to 15%. For some reviewers in our sample this conservatism bias goes up to 40%. Moreover, we find a significant, albeit small, effect of the reputation of reviewers on their strategies, indicating that incentives to manipulate demand in order to prevent reputation updating are present in this industry. We estimate an average probability of giving a good review despite a negative signal of 5%

[1] A previous version of this paper, featuring a reduced-form estimation of expert bias, provides some evidence of this phenomenon.

among reviewers concerned by potential conflicts of interest.

In order to provide evidence of the performance of our econometric model, we carry out an estimation of our model on a random subsample of our data set and conduct out-of-sample predictions. We find that the predicted distribution of reviews matches closely the one observed in the data. We also compare our model to two alternative specifications: one ruling out strategic biases and one assuming that all experts have the same ability and this ability is common knowledge. The first one corresponds to a story in which experts differ in their ability but are nonetheless truthful when sending their recommendations. The second one corresponds to a story in which experts are truthful because their payoff does not depend on how well informed they are. We find that our model performs better than these two in predicting the outcomes in the out-of-sample data, suggesting that strategic biases and career concerns do play an important role in movie reviewing.

We present the related literature in the next section. Theoretical predictions and our empirical strategy are presented in sections 2.3 and 2.4. We introduce our dataset in section 2.5 and our empirical analysis in section 2.6. Out-of-sample predictions are provided in section 2.7. Section 2.8 provides a short conclusion.

2.2 Related Literature

Technically, our estimation method follows the literature on the estimation of voting games by building on the technique developed by Iaryczower and Shum [28]. They estimate the political bias of supreme court judges who aim to take the right decision between rejecting or confirming the decision of the lower court. The main conceptual difference between our two approaches is that political biases are exogenous whereas our strategic bias is endogenous as it depends on the prior as well as the ability and reputation of the expert. They rely on a two-step approach in which biases and abilities are recovered from a first-step estimation of conditional voting probabilities. By contrast, our identification and estimation strategies rely on a partition of the set of priors in which we directly embed the structure of the theoretical model.

Assessing the importance of strategic bias in expert reviews is critical as they may have a non negligible role in the failure or success of new products, as shown in Reinstein and Snyder [45] or in Boatwright et al. [9]. Gentzkow and Shapiro [23] show that media outlets tend to conform their news coverage to their readership's taste in order to build their reputation as reliable news sources. Deviation from truthful revelation decreases with the verifiability of the state of the world and with competition. Their framework differs from ours in that the expert cannot influence the precision with which the state of the world is observed. Some empirical papers have tackled other potential sources of bias: DellaVigna and Hermle [15] conduct a detailed study of biases caused by conflicts of interest in the motion pictures industry and do not find any evidence of biased reviews when the reviewer works for the company producing the movie under review. Although we control for conflicts of interest in our estimation, this question is not central to our paper. Our structural approach yields estimates of biases caused by conflicts of interest that are significant at 11%. Although our results are not strongly significant, the difference with DellaVigna and Hermle [15] might be caused by methodological differences as they use a reduced form approach. Dobrescu et al. [16] study the impact of media concentration on reviews in the books industries. The authors also find that reviewers tend to give better reviews to products from the same media outlet, although they attribute this effect to similarities in tastes. In addition, they find some evidence that professional reviewers are less favourable to first time authors compared to consumer reviews. However, they only recover an aggregate effect, do not attribute this effect to strategic considerations and do not control for the possibility that professional and amateur reviewers grade differently. These papers provide atheoretical statistical evidence rather than structural estimations and do not quantify the biases.

2.3 A reputational cheap talk framework

In this section, we present the reputational cheap-talk game that we later estimate and characterize the truthful revelation set, which is at the core of our identification strategy. This set is defined as all priors supporting a truthful revelation equilibrium for experts of all abilities.

2.3.1 Setting of the game

We study a reputational cheap-talk game in which an expert receives a noisy private information (or signal) about the state of the world, and can communicate her information to her audience. In the following, we will assume that the state of the world represents the quality of an experience good while the expert's audience consists of the consumers who can potentially buy the product.

Figure 2.3.1 presents the timing of the game.

We assume that the quality, θ, is either bad, θ_0, or good, θ_1. We denote by Θ the quality space, $\Theta = \{\theta_0, \theta_1\}$. The private signal of the expert, s, is either low, s_0, or high, s_1. The signal space is denoted by $S = \{s_0, s_1\}$. The expert receives a low signal when she perceives the quality as bad and a high signal when she perceives the quality as good.

The expert and consumers share a common prior belief about the quality of the product. We denote by μ the common prior on the quality being high, $\mu = pr(\theta_1)$. The report of the expert takes the form of a recommendation about the product. The expert can send a recommendation at no cost. The expert cannot certify that her recommendation matches her private information nor can consumers verify that

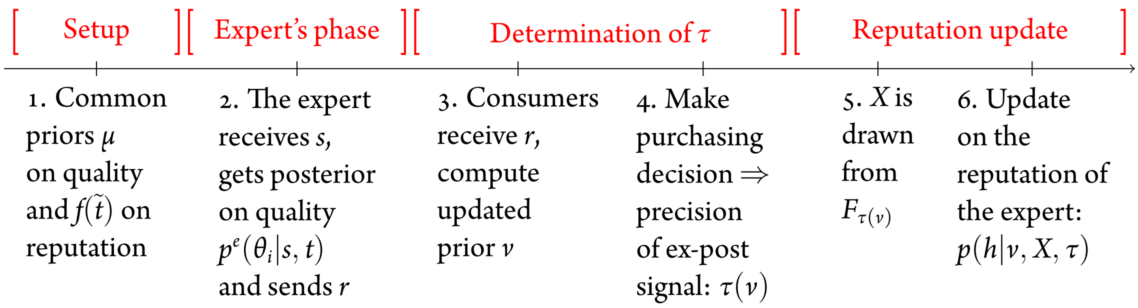

Figure 2.3.1: Timing of the game

the expert reported her private information truthfully. The recommendation of the expert, r, is either bad, r_0, or good, r_1. By sending a bad (good) recommendation the expert tells her audience that the movie is of bad (good) quality. The recommendation space is $R = \{r_0, r_1\}$.

The precision of the expert's private information depends on her ability. A more able expert receives a more precise signal. The ability t of the expert is defined as the probability for her to observe a private signal corresponding to the true quality, $p(s_i|\theta_i)$ in which $i = 0, 1$. t is privately observed by the expert. The expert's ability is drawn from a common knowledge distribution F over a continuum of possible abilities $\left[\frac{1}{2}, 1\right]$. The lowest ability expert receives a completely uninformative signal, $t = \frac{1}{2}$.

The expert cares only about her reputation, which is the consumers' belief about her ability, denoted $\tilde{t}$. The expert derives a utility $u(\tilde{t})$ of being perceived as having an ability $\tilde{t}$. We assume the expert's utility increases with her reputation. Furthermore, the expert is risk-neutral, therefore u is linear in $\tilde{t}$. Without loss of generality, we assume that $u(\tilde{t}) = \tilde{t}$. The expert's strategy, $\sigma : S \times R \rightarrow [0, 1]$, is the probability with which the expert sends the recommendation, r_i when she receives the signal, s_i: $\sigma_{s_i}(r_i) = pr(r_i|s_i)$.

Consumers can base their purchasing decisions on the expert's recommendation. They form an updated prior $\nu(\tilde{s}) = p(\theta_1|r, \tilde{t}, \tilde{\sigma})$ about the state of the world. The updated prior depends on the recommendation sent by the expert, the expert's reputation, and the consumers' belief about the strategy played by the expert, $\tilde{\sigma}$.

After consumption, each purchaser forms her individual opinion about the product's quality. The aggregation of the individual opinions of the purchasers forms an ex-post feedback that the entire consumer population use to compute their posterior belief about the quality. Regardless of their purchasing status, all consumers rely on the ex-post feedback derived from the aggregation of the purchasers' opinions to update their beliefs on the quality. Since everyone uses the same information, everyone shares the same beliefs. The precision of the ex-post feedback increases with the number of consumers who decided to buy the product and could therefore form their opinion about its quality. The precision of the ex-post feedback $\tau(.) \geq \frac{1}{2}$ is modelled as an increasing function of ν. A higher updated prior on the state of the world being high makes potential consumers more likely to buy. In turn, a higher number of purchasers translates into a more informative ex-post feedback. Formally, consumers observe an ex-post feedback signal $X \in \{X_0, X_1\}$ at the end of the game and $p(X_i|\theta_i) = \tau(v)$. Models with perfect ex-post revelation of the state of the world are equivalent to $\tau(v) = 1, \quad \forall v$.

Ultimately, consumers use their posterior belief on the quality to infer the probability with which the recommendation sent by the expert is correct and update the expert reputation accordingly.

In this game, the expert has two types of incentives not to reveal her private information truthfully.

(i) When the prior on the product's quality is extreme and the expert receives a signal that contradicts the prior, she anticipates that her signal is likely to be incorrect and has an incentive to lie and pretend having received the signal that is the most likely to be correct. We refer to that as *conservatism bias*.

(ii) Since consumers are less likely to buy when they expect the product's quality to be bad, by sending a bad recommendation, the expert can decrease the precision of the ex post feedback and prevent consumers from updating her reputation. Therefore, the expert has an incentive to over report the low signal. This incentive generates a *manipulation bias*.

The incentive for the expert to overreport the low signal to obfuscate the realization of the state of the world is stronger for highly reputed experts whose recommendation can change more drastically the consumers' updated prior on the quality and then influence more their purchasing decisions. Also simultaneous competition between experts mitigates the manipulation bias since a good recommendation written by a competitive expert reduces the deference effect on consumer demand that a bad recommendation can exert.

2.3.2 Truthful Revelation Set

We now characterize the truthful revelation set, which is the set of priors sustaining a truthtelling equilibrium for experts of all abilities.

The expert is said to be truthtelling if she sends a recommendation r_i after receiving a signal s_i with probability one, for $i = 0, 1$, that is $\sigma_{s_i}(r_i) = 1 \, \forall \, i \in 0, 1$. We denote σ^T the truthtelling strategy.

When consumers think that she is truthful, an expert has no incentives to deviate from truthtelling as long as she expects a higher utility for reporting truthfully her private signal than for misreporting it. Mathematically this condition writes as:

$$\mathbb{E}_f(u(\tilde{t})|s_i, r_i, t, \tilde{\sigma}^T) \geq \mathbb{E}_f(u(\tilde{t})|s_i, r_{-i}, t, \tilde{\sigma}^T), \qquad \forall i$$

Where $\tilde{\sigma}^T$ means that consumers believe the expert is truthtelling. In the following, to simplify notations, we sometimes do not mention the distribution on which we compute the expected utility, i.e. f. We also replace the pair $\{r_i, \tilde{\sigma}^T\}$ by $\tilde{s}_i$ which is the signal received by the expert as believed by consumers. When consumers receive r_i and believe that the expert is truthful, they think the expert has received s_i. (Their belief on the signal received by the expert is the information they use to update her reputation.) This leads us to the following characterization of the truthful revelation set. Its proof can be found in Appendix B.1 along the proofs of all propositions in this section:

Proposition 6. *Suppose consumers believe experts of all ability are truthtelling. After receiving s_i, $i \in \{0, 1\}$, the expert has no incentives to deviate from truthtelling if and only if she believes that she is more likely to be perceived by consumers as being right when she tells the truth ($r_i = s_i$) than when she lies ($r_i \neq s_i$):*

$$p^e(\tilde{\theta}_i | r_i, s_i, t, \mathbb{E}(\tilde{t})) \geq p^e(\tilde{\theta}_{-i} | r_{-i}, s_i, t, \mathbb{E}(\tilde{t})) \tag{2.3.1}$$

$\tilde{\theta}_i$ corresponds to the consumers' expectation on the quality. p^e denotes the probability computed by the expert.

Combining inequalities 2.3.1 for s_0 and s_1 yields a set of priors for which truthtelling is incentive compatible. This set of priors depends on the ability and the reputation of the expert. We denote this set $IC_{t,\mathbb{E}(\tilde{t})}$. As seen in the proof of proposition 6, only the expected ability of the expert affects the set of incentive compatible priors and not other moments of the distribution $f(\tilde{t})$. This makes the analysis considerably easier as we only need the first moment of the reputation to characterize the truthful revelation set.

Corollary 1. *In models with perfect revelation of the state of the world, $IC_{t,\mathbb{E}(\tilde{t})} = [1 - t, t]$.*

Corollary 1 shows that the reputation of the expert, $\mathbb{E}(\tilde{t})$, does not affect the set of incentive compatible priors when the state of the world is perfectly revealed ex-post. This set is also centred around $\frac{1}{2}$.

We also show that low ability experts have less incentives to be truthful:

Proposition 7. *Fix $\mathbb{E}(\tilde{t})$. Then, $IC_{t_1,\mathbb{E}(\tilde{t})} \subseteq IC_{t_2,\mathbb{E}(\tilde{t})} \Leftrightarrow t_1 \leq t_2$*

Hence, the set of priors for which truthtelling is incentive compatible for the least able expert, i.e. $IC_{\underline{t},\mathbb{E}(\tilde{t})}$, is included in all other sets. It is direct to see that for such priors, truthtelling is an equilibrium for all types of experts. We call this set the *truthful revelation set*.

Figure 2.3.2 gives the shape of such a set for different consumer beliefs about the ability of the expert when $\underline{t} = 0.65$.

In this paper, we do not characterize the equilibrium outside the truthful revelation set as it does not affect our estimation strategy.

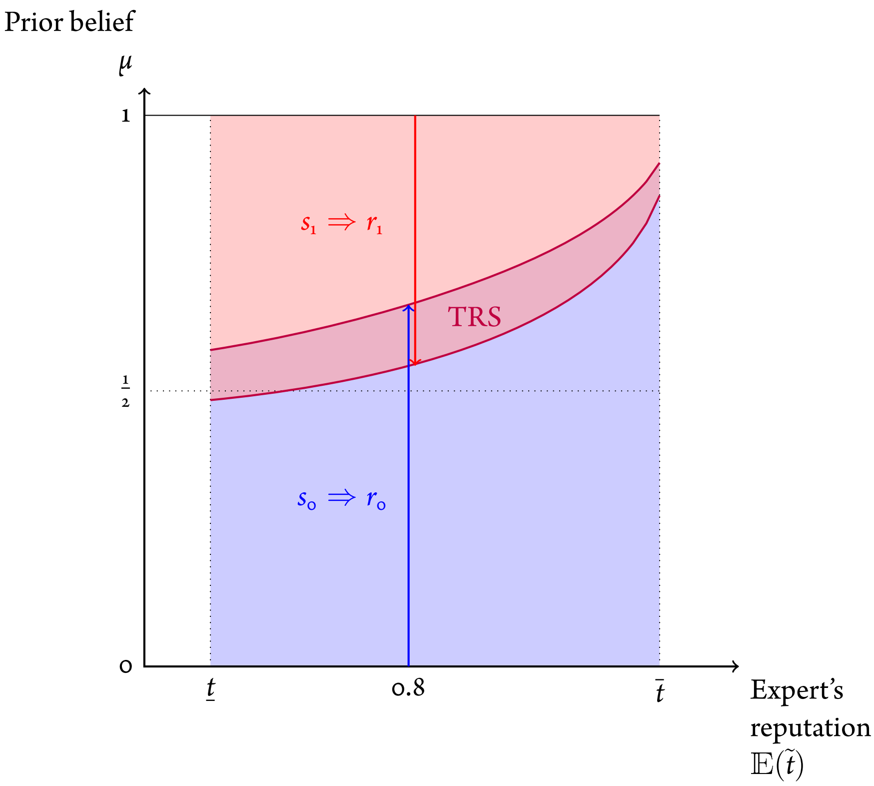

Figure 2.3.2: Truthful Revelation Set for $\underline{t} = 0.65$ and $\tau(\nu) = \frac{1+\nu}{2}$. The TRS lies between the two purple lines. These lines are determined by the set of incentive compatible priors for s_0, s_1, and all possible values of $\mathbb{E}(\tilde{t})$.

2.4 Estimation Strategy

2.4.1 Identification of Ability and Bias

Intuition

We are able to separately identify movies' priors, reviewers' abilities, and strategic biases. We use the common value in reviews to identify the priors: since reviews depend on reviewers' private signals and are released simultaneously, they are correlated through and only through the true quality. Indeed, a high quality movie is more likely to generate high private signals, and then good reviews than low quality ones. Holding the biases and abilities constant, the prior is identified as movies with extreme priors will tend to have more unanimous reviews. For a given prior, the partition on the set of priors created by the truthful revelation set allows us to identify each expert's ability and bias. We use observations with priors falling within the truthful revelation set to identify the abilities. For movies whose priors lie in this set, reviewers reveal truthfully their signals and their ability is given by the distribution of their reviews conditional on the quality. Outside the truthful revelation set, the strategic bias is estimated using the difference in the actual distribution of recommendations and the distribution of signals generated by the ability. Finally, the truthtelling set may vary according to the reputation of the expert. We identify the experts' reputation using this shift in the truthful revelation set.

A more formal argument

We now present a more formal argument for the identification. To obtain our result we make the following assumption on the game played by movie reviewers:

Assumption 1. *Experts play truthfully in the truthful revelation set.*

This assumption amounts to say that experts are truthful when truthtelling is an equilibrium.

Let us introduce some notations:

- $r_{i,j} \in \{0, 1\}$ is the recommendation or review given by the movie reviewer *i* to the movie *j* in our database. It corresponds to the recommendation in the theoretical model.
- $\mathbf{r}_j$ is the vector of reviews for movie *j*.
- $\mu_j = Pr(\theta_j = 1)$ is the prior belief about the movie. It is movie specific and common to all reviewers.[2]

[2] As in the theoretical model, we assume away heterogeneous priors. Identification does not hold in that case.

- $\gamma_{i,\theta} = Pr(r_{i,j} = 1|\theta_j = \theta, \mu_j)$ for $\theta = 0, 1$ are the reduced-form conditional probabilities of sending a good review.

The first step consists in maximizing the likelihood of observing the vector of reviews $\mathbf{r}_j$:

$$\max_{\{\gamma_{i,1},\gamma_{i,0}\}_{i=1}^{n},\{\mu_j\}_{j=1}^{m}} Pr(\mathbf{r}_j) = \mu_j \prod_{i=1}^{n} \gamma_{i,1}^{r_{i,j}}(1-\gamma_{i,1})^{1-r_{i,j}} + (1-\mu_j) \prod_{i=1}^{n} \gamma_{i,0}^{r_{i,j}}(1-\gamma_{i,0})^{1-r_{i,j}}$$

$$\text{s.t. } \gamma_{i,1} \geq \gamma_{i,0} \quad (2.4.1)$$

Since reviews are independent conditional on the quality, $\mathbf{r}_j$ follows a multivariate mixture distribution with mixing probabilities μ_j. Suppose first that the conditional probability of a good review, $\gamma_{i,\theta}$, is invariant over the movie prior. That would be the case if the strategic bias did not depend on the prior. In this setting, the nonparametric identification of the mixing probability and marginal distributions has been proven in several papers dealing with identification of mixture distributions such as Allman et al. [1] as long as we observe at least three reviews for each movie and we assume that $\gamma_{i,\theta_1} \geq \gamma_{i,\theta_0}$. The last inequality holds in our model since we assume that all reviewers have at least a 50% chance of receiving the correct signal about the quality.

However, in our framework, the conditional probability of giving a good review depends on the prior since the reviewer's strategy varies with the prior. We rely on an exclusion restriction in order to identify our structural parameters.

To precise how our identification strategy works, let us detail the reduced-form grading probabilities $\gamma_{i,\theta}$ according to our structural assumptions. In the following, we denote $\underline{\mu_i}$ and $\overline{\mu_i}$ the infimum and supremum of the truthful revelation set of movie reviewer i. Although it does not appear in the notations, we allow these bounds to depend on time in the estimation as reputation may vary. We also denote $b^- = p(r = 0|s_1)$ and $b^+ = p(r = 1|s_0)$ the expert's negative and positive bias.

Table 2.4.1: Probabilities of giving a good review conditional on the true state of the world and on the prior

Prior	$\gamma_{i,1}$	$\gamma_{i,0}$
$\mu_j < \underline{\mu_i}$	$t(1-b^-)$	$(1-t)(1-b^-)$
$\mu_j \in \left[\underline{\mu_i}, \overline{\mu_i}\right]$	t	$1-t$
$\mu_j > \overline{\mu_i}$	$t+(1-t)b^+$	$(1-t)+tb^+$

Note that the expert always reveals her private signal if this signal confirms the prior: the bias only concerns contradictory signals. Substituting the expressions in Table 2.4.1 in the likelihood function yields:

$$\max_{\{t_i, b_i^-, b_i^+, \underline{\mu}_i, \overline{\mu}_i\}_{i=1}^n, \{\mu_j\}_{j=1}^m} Pr(\mathbf{r}_j) =$$

$$\mu_j \prod_{i=1}^{n} \Big\{ 1\left(\mu_j < \underline{\mu}_i\right) \left[(t_i(1-b_i^-))^{r_i}(1-t_i(1-b_i^-))^{1-r_i}\right]$$
$$+1\left(\mu_j \in \left[\underline{\mu}_i, \overline{\mu}_i\right]\right)\left[t_i^{r_i}(1-t_i)^{1-r_i}\right]$$
$$+1\left(\mu_j > \overline{\mu}_i\right)\left[(t_i+(1-t_i)b_i^+)^{r_i}(1-t_i-(1-t_i)b_i^+)^{1-r_i}\right]\Big\}$$
$$+(1-\mu_j)\prod_{i=1}^{n}\Big\{ 1\left(\mu_j < \underline{\mu}_i\right)\left[((1-t_i)(1-b_i^-))^{r_i}(1-(1-t_i)(1-b_i^-))^{1-r_i}\right]$$
$$+1\left(\mu_j \in \left[\underline{\mu}_i, \overline{\mu}_i\right]\right)\left[(1-t_i)^{r_i}t_i^{1-r_i}\right]$$
$$+1\left(\mu_j > \overline{\mu}_i\right)\left[((1-t_i)+t_ib_i^+)^{r_i}(1-(1-t_i)-t_ib_i^+)^{1-r_i}\right]\Big\}$$

$$\text{s.t. } t_i \in \left[\frac{1}{2}, 1\right] \tag{2.4.2}$$

The intuition for identification clearly translates into the formulation of the likelihood. Our identification stems from this partition over the set of priors: the ability t_i is identified on $\left[\underline{\mu}_i, \overline{\mu}_i\right]$, whereas the negative and positive biases, b_i^- and b_i^+, are identified on $\left[0, \underline{\mu}_i\right]$ and $\left[\overline{\mu}_i, 1\right]$. Also, we can directly recover the parameters of interest of our model, t_i, b_i^-, and b_i^+, since they are related with the reduced-form conditional grading probabilities in a simple way.

2.4.2 Estimation

We allow the prior μ_j to depend parametrically on movie characteristics ω_j via the following logit formulation:

$$\mu(\omega_j;\beta) = \frac{\exp(\omega_j'\beta)}{1+exp(\omega_j'\beta)} \in [0,1]$$

Elements in ω_j include information on movies available to the market prior to their release, for instance the experience of the director measured as her number of previously directed movies, the production budget, a proxy for the genre of the movie such as MPAA ratings, whether or not the movie has been produced in the US, whether the movie is an original production, a remake, or a sequel.[3]

Similarly, we have to estimate the initial reputation of the reviewer as it will affect her truthtelling revelation set $\left[\underline{\mu}_i, \overline{\mu}_i\right]$. More precisely, as seen in the theoretical section, $\underline{\mu}_i$ and $\overline{\mu}_i$ are both functions $\underline{\varphi}$ and $\overline{\varphi}$ of the prior expectation of consumers about the expert's ability, $\mathbb{E}(t_i)$. We compute these functions using the characterization of $IC_{\underline{t},\mathbb{E}(\tilde{t})}$, provided by equation 2.3.1. We estimate $\mathbb{E}(t_i)$ via the following logit formulation:

$$\mathbb{E}(t_i|\omega_i;\delta) = \frac{\exp(\omega_i'\delta)}{1 + exp(\omega_i'\delta)} \in \left[\frac{1}{2}, 1\right]$$

in which ω_i is a shifter for the experts' reputation.

Since we need the reputation to be at least greater than one half, we impose δ to be greater than zero.[4] Since we choose ω_i to be a shifter for the expert's reputation, we consider this assumption as innocuous. Also note that we allow the reputation as well as ω_i to vary through time.

Also, to simplify notations, we introduce the following denominations for the partition over the set of possible priors:

- $\mathbf{I_A}(\omega_i;\delta) = \left[0, \underline{\varphi}(\mathbb{E}(t_i|\omega_i;\delta))\right]$: the set of priors lower than priors in the truthful revelation set. In this set, the expert can be affected by negative bias.
- $\mathbf{I_B}(\omega_i;\delta) = \left[\underline{\varphi}(\mathbb{E}(t_i|\omega_i;\delta)), \overline{\varphi}(\mathbb{E}(t_i|\omega_i;\delta))\right]$: the truthful revelation set.
- $\mathbf{I_C}(\omega_i;\delta) = [0, \overline{\varphi}(\mathbb{E}(t_i|\omega_i;\delta))]$: the set of priors greater than priors in the truthful revelation set. In this set, the expert can be affected by positive bias.

All these sets depend on the reputation of each expert via δ and ω_i.

$\left\{t_i, b_i^-, b_i^+\right\}_{i=1}^n$ are reviewer-specific and estimated using only the distribution of reviews.

All the parameters are estimated by maximizing the following likelihood function:

[3] Including the actors would be difficult as an actor's notoriety can vary greatly during her career and we would therefore need the actors' reputation for all movies in our panel.

[4] Remember that the expert's reputation is the expectation on her ability, and $\underline{t} > \frac{1}{2}$.

$$
\max_{\left\{t_i, b_i^-, b_i^+\right\}_{i=1}^n, \beta, \delta}
$$

$$
\begin{aligned}
\sum_j \log \Bigg[\mu(\omega_j;\beta) \prod_{i=1}^n \Big\{ & 1\Big(\mu(\omega_j;\beta) \in \mathbf{I_A}(\omega_i;\delta)\Big) \Big[(t_i(1-b_i^-))^{r_i}(1-t_i+t_ib_i^-)^{1-r_i}\Big] \\
& +1\Big(\mu(\omega_j;\beta) \in \mathbf{I_B}(\omega_i;\delta)\Big) \Big[t_i^{r_i}(1-t_i)^{1-r_i}\Big] \\
& +1\Big(\mu(\omega_j;\beta) \in \mathbf{I_C}(\omega_i;\delta)\Big) \Big[(t_i+(1-t_i)b_i^+)^{r_i}((1-t_i)(1-b_i^+))^{1-r_i}\Big] \Big\} \\
+(1-\mu(\omega_j;\beta)) \prod_{i=1}^n \Big\{ & 1\Big(\mu(\omega_j;\beta) \in \mathbf{I_A}(\omega_i;\delta)\Big) \Big[((1-t_i)(1-b_i^-))^{r_i}(t_i+(1-t_i)b_i^-)^{1-r_i}\Big] \\
& +1\Big(\mu(\omega_j;\beta) \in \mathbf{I_B}(\omega_i;\delta)\Big) \Big[(1-t_i)^{r_i}t_i^{1-r_i}\Big] \\
& +1\Big(\mu(\omega_j;\beta) \in \mathbf{I_C}(\omega_i;\delta)\Big) \Big[((1-t_i)+t_ib_i^+)^{r_i}(t_i(1-b_i^+))^{1-r_i}\Big] \Big\} \Bigg]
\end{aligned}
$$

$$
\text{s.t. } t_i \in \left[\frac{1}{2}, 1\right], \quad b_i^+, b_i^- \in [0,1], \delta \geq 0 \tag{2.4.3}
$$

We provide a more robust estimation strategy taking into account potential conflicts of interest between the reviewer and the production company of the movie in Appendix B.2.

We conclude this section by stating two assumptions we used in our estimation:

Assumption 2. $\underline{t} = 0.55$

Assumption 3. *The precision of the aggregate signal X is linear in the intermediate posterior of the consumers, more precisely:* $\tau(\nu) = \frac{1+\nu}{2}$.

By definition, the truthful revelation set is the set of incentive compatible priors corresponding to the lowest ability $\underline{t} \geq \frac{1}{2}$. Since we do not know the value of the lowest ability in reality, a robust approach would be to choose $\underline{t}$ as close to one half as possible. Indeed, by proposition 7, the chosen truthful revelation set would lie within the actual one. However, the truthful revelation set converges to a line when $\underline{t}$ tends toward one half, making it impossible to identify the ability.[5] We therefore choose $\underline{t} = 0.55$ in our estimations. We feel this value is close enough to one half to be robust, a feeling confirmed by our results: all our estimates of the reviewers' ability are well above this threshold.

[5]When the truthful revelation set is a line, the probability that an observation falls within is zero.

Assumption 3 is a simplifying assumption restricting the way consumers react to the reviews. In the future, we plan to estimate a parametric formulation of $\tau(v) = \frac{1+v^a}{2}$ with $a \in \mathbb{R}^+$.

A graph depicting the truthful revelation set used for our estimations can be found in Appendix B.4. The shape of this set depends on these two previous assumptions.

2.5 Presentation of the Data

We use movie ratings by professional movie reviewers published in the rottentomatoes.com website. This website gathers reviews from the most prominent movie reviewers in North America. The reviews provided by rottentomatoes correspond to rescaled grades of the original reviews' numerical or alphabetical scores. When the original review is associated with no grade, or the grade is average, for instance 3 stars out of 5, rottentomatoes assigns a rating based on the tone of the review. Hence, each review is associated with a *fresh* or *rotten* rating, a binary structure completely in line with our theoretical model and structural technique of estimation.

We restrict our dataset to ratings by reviewers qualified as *top critics* by rottentomatoes on movies released in the US between 1990 and 2013. The reviewers in our sample are therefore professional reviewers likely to be driven by career concerns, which is in line with our theoretical model. The dataset has a panel structure whose two dimensions are the movies and the reviewers.

Each review contains the newspaper or TV show in which it was given, the name of the reviewer, and the date of the publication. We only keep reviews which are close in time so as to ensure that competition is indeed simultaneous and no herding behaviour occurs.

We combine these reviews with data on movies' characteristics, collected on the imdb.com website, i.e. the *Internet Movie DataBase*. For each movie, the dataset includes the official release date, the total box-office revenues, the production budget, the number of screens on which the movie was run during the opening weekend, the MPAA rating, the genre and whether or not the movie was produced in the US. In addition, for each movie, the dataset contains the director's name and the number of films she directed in the past.

Finally, we include some reviewers' characteristics to control for their reputation, namely an index summarizing the number of monthly searches on Google provided by the website Google Trends, and some characteristics of the newspapers or TV shows for which they gave their reviews. We also control for potential conflicts of interest by identifying cases in which the newspaper and the production company of the movie belong to the same media outlet.

Our complete sample contains a total of 118,208 reviews over 5,578 movies by 1,242 reviewers. The variables in our dataset are extensively detailed in Appendix B.3.

In the following, we describe more precisely how we build our dataset.

2.5.1 Movie reviewers

Rottentomatoes defines its *top critic* category on the basis of the size of their audience.[6]

We focus on the more prolific reviewers in our sample, i.e. those who individually published more than 850 reviews.[7] Indeed, we need to observe many reviews per reviewer to be able to estimate their individual abilities. 44 of them individually published more than 850 reviews, our most prolific reviewer being the famous Roger Ebert with 3,277 reviews.

We choose the reputation shifter of the reviewers ω_i to be an index representing the number of monthly Google searches for each of them. This index is provided by the website Google Trends and takes values between 0 and 100, 100 being attributed to the point in time when the popularity of the reviewer is maximal.[8]

To take into account the experts' reputation, we have to disregard nine reviewers as they have highly popular homonyms and we cannot disentangle their reputation from the reputation of the celebrity they share their name with. For instance, we have to exclude a Roger Moore, who clearly never played the role of a famous British secret service agent.[9] We end up with 35 reviewers, whose list is found in Appendix B.3. Overall, these reviewers produced 48,278 reviews, which represents about 41% of our total number of observations.

We also observe the medium through which the reviewer gave her review and the precise date of the review. Reviewers in our sample publish in various kinds of media, from national TV and radio shows to local newspapers. We are able to identify whether or not the medium is addressed to a general or a specialized audience, whether or not its coverage is nationwide or statewide. The date of the review allows us to ensure that we consider a simultaneous competition game. Since Google Trends data only start in 2004, we only consider reviews given from 2004 to 2013. We also exclude movies released prior to 2003

[6]According to Rottentomatoe, "To be considered for Top Critics designation, a critic must be published at a print publication in the top 10% of circulation, employed as a film critic at a national broadcast outlet for no less than five years, or employed as a film critic for an editorial-based website with over 1.5 million monthly unique visitors for a minimum of three years. A Top Critic may also be recognized as such based on their influence, reach, reputation, and/or quality of writing, as determined by Rotten Tomatoes staff."

[7]The threshold of 850 reviews is arbitrary: it gives the maximum number of movie reviewers for an acceptable amount of reviews.

[8]When comparing the popularity of two reviewers, the returned values are computed relative to the maximum number of searches for either of them over the period. Since the most popular reviewer in our sample, Roger Ebert, is clearly an outlier in terms of reputation and we do want to observe some variations between reviewers' reputation, we choose a moderately popular reviewer, Peter Travers, as our benchmark and compare all the others to him.

[9]Our systematic rule to determine whether or not a reviewer has a famous homonym is typing their name in a web search engine and checking that all first results concern a movie reviewer.

as some movies of the 1990s are given late reviews for their DVD release ten years later.[10] We consider that these reviews could be biased by some herding behaviours which do not enter into the scope of our analysis. Our final dataset shows an average standard deviation of 39 days for the dates of the reviews. We aim to reduce this standard deviation by excluding late reviews in future estimations.

We finally exclude all movies with less than four reviews because: (a) we need variations in the reviews to identify the prior of the movie, (b) we want movies with a comparable number of reviews to lift the issue of a potential selection bias from reviewers. Our final sample contains 30,531 reviews over 2,413 movies by 35 reviewers.

2.5.2 Description of the Movies

Our dataset consists of movies released in the US between 1990 and 2013 which have been granted a rating by the Motion Picture Association of America. The MPAA ratings are guidelines given to parents on the contents of movies: they consist in the five following different grades G, PG, PG-13, R, and NC-17. Although these ratings are not central to our analysis, we choose to focus on MPAA rated movies because they represent a large part of all movies released in the US. Indeed, the six main production companies are part of the MPAA and must therefore submit their products to the rating system.

Table 2.5.1 presents some descriptive statistics on the movies in our sample. As shown by the converted budget in today's US$ and the width of the release as the number of screens in the opening weekend, our sample features blockbusters as well as more independent movies. Indeed, the average production budget for a wide release in the USA is $66 million which is twice our sample average.[11] Also, our median movie is released on only 179 screens in the opening week, which is below the cut-off value above which you can consider a movie as a wide release: 600 screens according to Einav [18]. That does not necessarily mean that lots of movies in our sample are small confidential productions, also known as movies in limited release, but more probably some of our movies have been first released in a few cities to build a reputation and then have been displayed on more screens. These movies are called *platform release*. The US gross profit expressed in today's US$ also shows a diversity in the way movies have succeeded, from big hits to big failures, with a relatively high standard deviation. # weeks stands for the total number of weeks during which the movie has been screened on some theater in the US. On average a theater runs a movie from 6 to 8 weeks, but some might run it longer, so our median total number of 11 weeks seems consistent with this observation. (See Einav [18] for figures on the average run of a movie.) We also include the number of previous movies by the same director, which should influence the prior on the movie, and is concentrated around low values in our dataset, with a long thin tail.

[10] We choose to include 2003 to keep the movies released in late 2003 and reviewed in early 2004.
[11] Source: www.the-numbers.com

Table 2.5.1: Some descriptive statistics on the movies

	B.O. US (M$)	# weeks	# screens 1st week	Budget (M$)	#Previous Movies
# obs.	4821	4563	4312	3556	5578
Median	9.7	11	179	23	3
Mean	38.9	12.15	1195	39	5.33
St. Dev.	68.7	9.32	1342	45	7.18
Highest	Titanic: 951	Roving Mars: 167	The Dark Knight Rises: 4404	Pirates of the Caribbean 3: 335	Chunhyangdyun (Im, Kwon-taek): 96
2nd	Avatar: 821	Deep Sea: 165	Iron Man 2: 4380	Titanic: 289	Éloge de l'amour (Godard, Jean-Luc): 74
3rd	The Avengers: 629	Aliens of the Deep: 142	Harry Potter and the Deathly Hallows - Part 2: 4375	Spider-Man 3: 288	Kakushi ken oni no tsume (Yamada, Yôji): 72

2.5.3 Media Outlets

We conclude the description of our data by detailing the potential conflicts of interest we find in our sample. We follow DellaVigna and Hermle [15] and create a dummy variable equal to 1 if the production company of the movie and the medium of the review belong to the same media outlet at the time of the review. Since this question is not central to our paper but included to show that our results are not driven by the wrong kind of bias, we do not provide a description as detailed as in their paper.

We identify a new potential source of conflicts of interest: the Disney Media Group which controls the Ebert & Roeper TV show and studios such as Miramax and Walt Disney Pictures.

Overall, only 1.4% of our observations are affected with these conflicts of interest. All possible conflicts of interest that we are aware of in our sample are detailed in Appendix B.3. Note that when reviewers express themselves on several media, we only consider the reviews in media in the same group as the production companies to be problematic.

2.6 Results

We now turn to the results of our empirical analysis. Table 2.6.1 and Table 2.6.2 present the movie-specific estimates and a summary of the reviewer-specific estimates under two alternative specifications. Compared to specification (I), specification (II) includes the budget in the determination of the prior and takes into account potential conflicts of interest. Table B.4.1 in Appendix B.4 gives the exhaustive list of the reviewer-specific parameters' estimates.

Table 2.6.1: ML estimates for the prior, reputation, and conflicts of interest — Specification (II) includes the budget and conflicts of interest

	(I)		(II)	
	Coeff.	Bootstrap SE	Coeff	Bootstrap SE
Movie Specific, β:				
Constant	1.249	(0.040)	1.124	(0.004)
Origin: USA	-1.494	(0.038)	-1.440	(0.016)
Origin: co-production USA	-1.539	(0.074)	-1.554	(0.083)
Remake	-0.501	(0.247)	-0.439	(0.228)
Sequel	-0.496	(0.053)	-0.432	(0.115)
Number of director's previous films	0.016	(0.000)	0.016	(0.000)
G rating	0.949	(0.173)	0.988	(0.332)
PG rating	0.087	(0.076)	0.470	(0.191)
R rating	0.484	(0.043)	0.530	(0.040)
NC-17 rating	0.690	(9.160)	0.650	(4.117)
log budget			0.000	(0.000)
Reputation, δ:				
Google Search Index	5.9×10^{-4}	(0.000)	5.9×10^{-4}	(0.000)
Conflict of interest, b^c:				
Average Bias			0.055	(0.034)
# Observations:	30440		22674	

Notes: Bootstrap Standard Errors are computed on 100 iterations for (I) and 500 for (II)

The prior depends positively and significantly on a movie being a foreign production. Our interpretation is that low quality foreign films are in general not successful enough to be released in the US. When facing a foreign movie, consumers rationally expect it to be of high quality. Remakes and sequels also impact negatively the prior. The result on sequels is in line with an empirical observation that consumers give significantly lower grades to sequels: in general the sequel must be of lower quality than the original work, hence the lower prior. The same story can apply to remakes, although the coefficient in this case is less significant. As expected, the director's experience impacts positively and significantly the prior on the movie. Compared to a PG-13 MPAA rating, G and R ratings impact positively the prior whereas PG and NC-17 are not significant. We use these ratings as controls for the content of the movies. Surprisingly, the budget has no strong impact on the prior. Our interpretation is that the gains from a larger budget, i.e. better actors, special effects, etc., might be counterbalanced by a loss in originality and quality from big productions.

The coefficient on the Google search index is quite small and suggests that consumers' manipulation by movie reviewers is limited. With such a coefficient, only Roger Ebert would have had a real shot at decreasing demand to maintain his reputation, but as the estimates of his negative bias suggests, he did not use this mechanism. However, this result means the model with endogenous realization of the state of the world is a better fit to the data than the model in which we learn the state of the world at the end of the game. Indeed, the estimates do suggest that the truthful revelation set is shifted towards high prior movies if the reputation is high enough. If experts were not able to manipulate demand, $\hat{\delta}$ would have been equal to zero.

The small amplitude of the effect might be an artifact of two assumptions: (a) the linearity of the precision of the aggregate signal, (b) the implicit assumption that movie reviewers operate in autarkic markets. To check the robustness of our findings, we can (a) estimate the precision of the aggregate signal as a function of the prior and the reviews, (b) estimate a model in which movie reviewing is part of a competitive market.[12]

Finally, we find a positive bias caused by some conflict of interest. This bias is quite large since reviewers turn a bad review into a good one 5% of the time when reviewing a movie from the same media outlet. It is significant at an 11% level.

All our estimates are robust to the change in specification from (I) to (II). This is in line with the fact that the budget has no effect on the prior and that conflicts of interest should also be independent from the movie prior.

Table 2.6.2 summarizes our findings on the reviewer-specific parameters, i.e. their ability, negative bias, and positive bias. First, all the estimates are robust to the change in specifications: the number of instances in which we have a conflict of interest is probably too small to bias the estimates of b^- and b^+ in (I). The average ability of experts in our sample is quite high, with movie reviewers being able to correctly observe the quality of a movie 78% of the time. The average positive and negative biases are significant with reviewers reporting a negative opinion after receiving a positive signal and vice versa 8% and 15% of

[12] In that case, the fact that we exclude movies with less than four reviews necessarily leads to an underestimation of the feasibility of manipulation.

Table 2.6.2: Summary of ML estimates for reviewer-specific parameters — Specification (II) includes the budget and conflicts of interest

		(I)			(II)		
		t	b^{-}	b^{+}	t	b^{-}	b^{+}
Median		0.78	0.04	0.11	0.77	0.08	0.11
Mean		0.77	0.08	0.15	0.77	0.011	0.13
St. Dev.		0.06	0.11	0.13	0.06	0.011	0.13
Highest ability	Robert Denerstein:	0.898	0.107	0.001	0.898	0.146	0.013
		(0.009)	(0.035)	(0.034)	(0.021)	(0.062)	(0.042)
Lowest ability	Kyle Smith:	0.619	0.345	0.000	0.654	0.363	0.000
		(0.015)	(0.047)	(0.000)	(0.026)	(0.065)	(0.000)
# Observations:			30440			22674	

Notes: Bootstrap Standard Errors are computed on 100 iterations for (I) and 500 for (II)

the time.

The highest ability reviewer, Robert Denerstein, has a precision of 90% and no bias on high prior movies. However, for low prior movies, he will send a bad review after observing a positive signal 11% of the time. Kyle Smith, the lowest ability reviewer, only observe the accurate signal 62% of the time and has a strong negative bias of 34%.

A surprising feature of our results is that some reviewers are apparently strongly polarized, i.e. they are either prone to negative criticism or to positive criticism. For instance, Liam Lacey has a strong negative bias of 41%. But he is truthtelling for movies with positive priors. On the opposite, Joe Baltake is completely truthtelling for low prior movies, but has a strong positive bias of 43% for high prior movies. At this point, we do not know whether these differences are mere features of the strategies played at equilibrium, if they are explained by different personal tastes towards negative or positive criticism, or if they indicate that reviewers are heterogeneous in their abilities to recognize high and low qualities.

Finally, it does seem that more able experts are less prone to bias: we find a slightly negative relationship between bias and ability. However, having observations on only 35 reviewers, this relationship is not significant.

2.7 Out-of-sample Predictions

In this section, we test the predictive power of our model. We find that the predicted distribution of reviews matches closely the distribution of reviews observed in the data. Our model also better fits out-of-sample data than two alternative models: one ruling out strategic biases and one assuming that all experts are equally well informed and that their ability is common knowledge.

2.7.1 Testing Predictive Power

Applying our model's estimates to new data, i.e. new movies, we can recover the priors on their quality but also the ex-ante probability, i.e. prior to observing the private signal, that a particular reviewer gives a good or bad review to a movie. However, we only observe the realization of this ex-ante distribution, which is the final review. Testing the predictive power of our model is therefore difficult since we cannot directly compare the estimated and realized outcomes.

To overcome this difficulty, we group together observations, i.e. couples of movie and movie reviewer, for which our model predicts similar probabilities of good reviews. Within each group, we then compute the actual proportion of good reviews which corresponds to the observed probability of a good review conditional on belonging to the group. With this method, we cannot generally conclude that a model with similar predicted and observed probabilities is a good predictor.[13] We can however rule out models whose predicted probabilities differ widely from the actual ones.

We provide an additional proof of the relevance of our model by comparing its predictive power to the one of alternative models. To achieve this goal we compute the confusion matrix of our model: for each movie-movie reviewer couple in our out-of-sample data, we draw a realization of the review from our estimated ex-ante distribution of giving a good or a bad review. We then compare this realization to the actual review given by the reviewer. If they are similar, we say our model classified the observation correctly. Repeating that several times, we can obtain an average classification ratio. Doing that same operation for alternative models, the best one has the highest average classification ratio.

Out-of-sample Data

We randomly split our main data set used in section 2.6 and estimate our model on 80% of the movies in our sample. The 20% remaining, which form our out-of-sample data set, are used to test the out-of-sample predictive power of our model. Out of the 1661 movies and 22674 reviews in the main data set, the out-of-sample data set contains 351 movies for a total of 4792 reviews. This out-of-sample dataset needs to be representative. This is in theory ensured by the fact that it is constructed randomly. As additional evidence, Table 2.7.1 shows that moments of several variables are similar between our main data set and the out-of-sample data set.

[13] Especially if the rule allocating an observation to a group is too coarse. For instance, if we take all observations in one group, the average review in the main data is a pretty good estimator of the average review in the out-of-sample data.

Table 2.7.1: Comparison of Moments between the Main Data Set and the Data Set Used for Out-of-sample Predictions

	Main	Out-of-sample
Avg. Budget (M$)	43.7	41.9
Director's Avg. Number of Previous Movies	6.50	6.95
Freq. of US productions	57.3%	57.0 %
Avg. Number of Reviews	13.65	13.65
Freq. of Good Reviews	49%	48%
Proportion of movies released in:		
2003	2.1%	4.3%
2004	11.3%	12.0%
2005	12.5%	10.3%
2006	12.3%	12.3%
2007	11.6%	10.5%
2008	11.4%	11.7%
2009	10.7%	10.3%
2010	9.3%	10.8%
2011	10.3%	8.8%
2012	7.2%	7.4%
2013	1.4%	1.7%

Model Specifications

Our main econometric model is model (II) from section 2.6. It includes the budget in the prior of the movies and potential conflicts of interest which can arise when the movie's production company and the reviewer's newspaper or show belong to the same production company. To carry out out-of-sample predictions, we need to estimate this model again using our restricted data set. Estimates are reported in Tables B.4.2 and B.4.3 of Appendix B.4 and are remarkably close to the ones using the full data set.

We compare this model to alternative theories. The first one, called model (III), assumes that reviewers have a private ability and observe a private signal but are always truthful. Model (III) is therefore a model excluding strategic biases. We estimate the parameters of this model by maximizing the following likelihood:

$$\max_{\{t_i\}_{i=1}^{n},\beta} \quad \sum_j \log\left[\mu(\omega_j;\beta)\prod_{i=1}^{n}\left[t_i^{r_i}(1-t_i)^{1-r_i}\right] + (1-\mu(\omega_j;\beta))\prod_{i=1}^{n}\left[(1-t_i)^{r_i}t_i^{1-r_i}\right]\right] \text{ s.t. } t_i \in \left[\frac{1}{2},1\right]$$

The second one, called model (IV), assumes that reviewers observe a private signal which depends solely on the publicly observable prior plus an i.i.d. noise, which is independent from their ability. In this model, reviewers also reveal truthfully their signal since their pay-off does not depend on their reputation for being well informed. Consequently, reviews follow the prior distribution, which we can estimate using a simple logit approach.

Estimates for the two alternative models are reported in Tables B.4.2 and B.4.3 of Appendix B.4.

2.7.2 Out-of-sample Fit and Classification Power

Our first approach consists in computing the ex-ante distribution of good and bad reviews as predicted by our model (II) for each observation in our out-of-sample data and grouping together the observations for which our model predicts a similar distribution of reviews. We then compare the predicted probabilities of good and bad reviews with the frequencies observed in each groups. Table 2.7.2 and Figure 2.7.1 show that our model performs quite well according to this criterion. In Table 2.7.2, we group predicted probabilities of good reviews by increments of 10%. For each group with a non-negligible amount of observations except one, the differences between the predicted and observed probabilities of a good review are less than 5%. The only exception arises in [0.5, 0.6], in which the difference is around 8%.

In Figure 2.7.1, out-of-sample observations are sorted horizontally according to their predicted probabilities of good reviews. The red curve represents a non-parametric fit of the actual proportion of good reviews. Groups here are defined by the kernel of the fit and the closer the red curve is to the

Table 2.7.2: Predicted Probabilities of Good Reviews and their Observed Frequency in the Data

$\widehat{P(r_{i,j} = 1)}$ $\in$	# Obser.	Avg. Predicted Probability of a Good Review	Observed Freq. of Good Reviews
[0, 0.1]	0	—	—
[0.1, 0.2]	1	18.73%	0
[0.2, 0.3]	88	26.37%	29.55%
[0.3, 0.4]	678	35.80%	39.38%
[0.4, 0.5]	1319	44.70%	47.38%
[0.5, 0.6]	1194	55.27%	47.32%
[0.6, 0.7]	1112	64.91%	66.19%
[0.7, 0.8]	356	74.70%	76.97%
[0.8, 0.9]	44	81.40%	77.27%
[0.9, 1]	0	—	—

45-degree line, the better our model is at predicting reviews.

We now compare confusion matrices for our model and for alternative models, namely (III) and (IV). Elements of the matrix are the number of observations corresponding to the predicted and observed outcomes. The overall classification ratio is similarly defined as the number of observations properly predicted over the total number of observations. A high classification ratio is indicative of a high predictive power. Figure 2.7.2 and Table 2.7.3 present the confusion matrices and classification ratios averaged over 100 repetitions for the three models under study.

Table 2.7.3: Classification Ratios Averaged over 100 Repetitions for Models (II), (III), and (IV)

	(II)	(III)	(IV)
Overall Classification Ratio	0.531	0.515	0.519
Classification Ratio on Good Reviews	0.558	0.513	0.555
Classification Ratio on Bad Reviews	0.501	0.518	0.479

Classification ratios are generally quite low. This is to be expected given that for a large part of our observations, our models predict probabilities of good reviews around one half. For this type of movies, the model misclassifies an observation one half of the time. This tool is therefore relevant only insofar as it allows us to compare several models. Despite some issues with the classification of bad reviews, with an

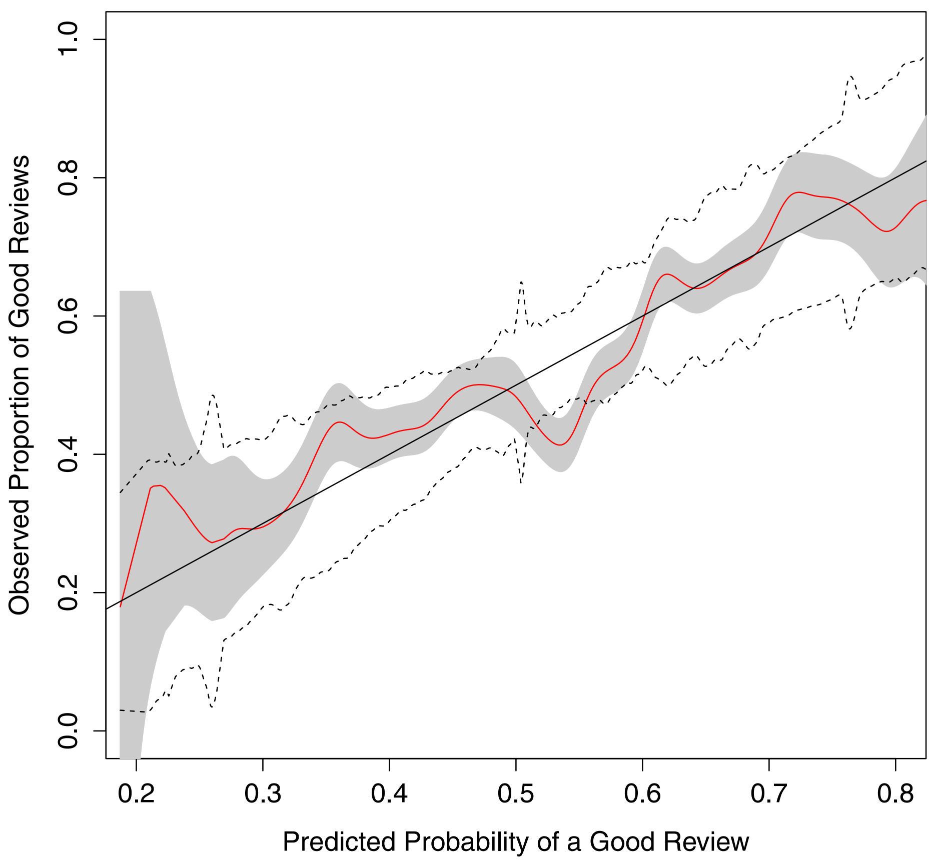

Figure 2.7.1: Observed Proportion of Good Reviews given their Predicted Probability — The red line displays the non-parametric fit of the observed proportion of good reviews. The shaded area represents the 95% confidence interval of this fit. The black plain line is the 45 degree line. The black dotted line represents the 95% confidence interval for the predicted distribution.

$$
\text{Predicted}
$$

$$
\text{Observed}\quad
\begin{array}{c} \\ 0 \\ 1 \end{array}\!\!
\begin{array}{cc} 0 & 1 \\ \hline \end{array}
$$

$$
\text{Observed}\quad
\underset{\text{Model (II)}}{\begin{array}{c|cc} & 0 & 1 \\ \hline 0 & 1134.5 & 1130.5 \\ 1 & 1117 & 1410 \end{array}}
\qquad
\underset{\text{Model (III)}}{\begin{array}{c|cc} & 0 & 1 \\ \hline 0 & 1174.2 & 1090.8 \\ 1 & 1231.8 & 1295.2 \end{array}}
\qquad
\underset{\text{Model (IV)}}{\begin{array}{c|cc} & 0 & 1 \\ \hline 0 & 1085 & 1180 \\ 1 & 1124.1 & 1402.9 \end{array}}
$$

Figure 2.7.2: Confusion Matrices Averaged over 100 Repetitions for Models (II), (III), and (IV) — The first row and column of each matrix represent observed and predicted bad reviews. The second row and column are similarly defined for good reviews.

overall classification ratio of 0.53, our model performs better than the alternative specifications. In addition to the good fit of predicted probabilities of good reviews to observed ones, this is a good indicator that models of expertise with reputational concerns and strategic biases best describe the behavior of movie reviewers.

2.8 Conclusion

We estimate the strategic incentives of experts to send biased recommendations using movie reviews. In our model, experts want to maximize their reputation as good predictors of the state of the world. Bias in this case can take two forms: experts can disregard contradictory signals because they are noisy, and they can send negative recommendations to discourage demand and hinder the update on their reputation. Expert bias depends both on the prior on the state of the world and on the expert through her ability and reputation.

To tackle these issues, and the fact that the state of the world, abilities, and private signals are not observable, we introduce new identification and estimation strategies. The prior is identified by the fact that signals, and therefore reviews, are correlated through the unobservable true quality. The experts' abilities are identified by the fact that experts are unbiased on a subset of the priors. Outside this subset, the difference between the actual distribution of the expert's reviews and the one generated by a truthtelling expert with the same ability gives us the bias. The estimation is a straightforward one-step process, which allows us to recover the determinants of the prior, the relative impact of the reputation, any bias caused by a conflict of interest, and the ability and biases of each expert in our sample.

We find strong variations in movie reviewers' abilities and biases. The average negative and positive biases in our sample are strong and significant, with information manipulation occurring on average 10% of the time. Also, we find a small but significant impact of reputation validating models with endogenous realization of the state of the world at the expense of models in which it is revealed.

We devote to future research the introduction of heterogeneous abilities to recognize high quality and low quality movies, the case of competitive review markets, the estimation of the precision of the aggregate signal, and a fully structural estimation of the reputation.

3

Market characteristics and implicit contracts when consumers have switching costs

co-authored with Guillem Roig

Abstract

We study the incentives for firms to engage in tacit collusion when consumer have switching costs. In addition to the level of consumer's switching costs and players' patience, market characteristics, such as its granularity or the consumers' sophistication and lifespan greatly impact the implementability of implicit contracts. With long-run consumers, collusion is harder to implement than with short-run ones, which indicates that durable goods markets are more likely to harbour collusive behaviours. However, this result does not hold when consumers cease to be sophisticated or when we switch to a continuum of consumers. When the level of switching costs increases with the number of periods a consumer have remained captive to a firm, collusion is harder to implement.

3.1 Introduction

We develop a duopolistic model in which an homogeneous product is sold to a group of consumers over an infinite number of periods. At the beginning of the game, consumers are captive to either firm and endure a switching cost if they buy from the competitor. We obtain that the implementability of tacit collusion depends on the type of consumers in the market. Figure 3.1.1 summarizes the results on the

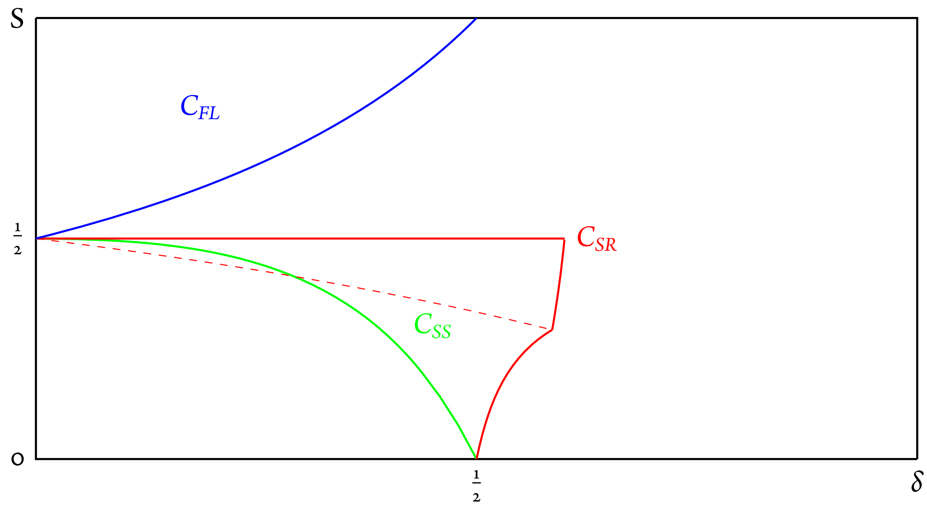

Figure 3.1.1: Implementability of full collusion given consumers' sophistication and lifespan (2 consumers) — Each plain line represents the lower boundary of the set over which full collusion is implementable, for a given consumer specification. The blue line stands for forward-looking consumers, the green line for short-sighted ones, and the red line for short-run ones (the dashed line represents an alternative equilibrium when consumers are short-run).

implementability of tacit collusion when there are two representative consumers in the market. In accordance to general wisdom, the discount factor positively affects the feasibility of tacit collusion. This happens whenever consumers are short-run — they interact with the firm only once — and long-run but short-sighted — they interact infinitely with the firm but are not able to anticipate the firms' future prices.[1] Conversely, a high discount factor makes it impossible to implement implicit contracts when consumers are long-run and forward-looking — those interact infinitely and are able to anticipate firms' future prices. In our model, competition is more severe with an asymmetric market structure and sophisticated consumers are more prompt to break a collusive outcome.[2] Moreover, collusion is harder to

[1]A short-run consumer with switching cost is for instance a tourist with a transportation cost. Alternatively, short-run consumers can be used here to model markets for durable goods.

[2]Therefore, the traditional correlation between indices of concentration (Herfindahl index) and industry profitability found in most competition models does not necessarily apply in models where consumers have switching costs such as the one studied

sustain as firm's deviation prices incorporate rents that consumers obtain after switching, and these rents are increasing with the weight consumers put on the future. Therefore, consumer sophistication hurts firms as it hinders the sustainability of implicit contracts.

The case of sophisticated consumers introduces many challenges: consumers are strategic and can coordinate to break collusion, but deviations of firms also depend on the consumers' expectation about their future utility when switching. Hence, in each period, after prices are posted, a coordination game is played among consumers who must decide between switching or staying loyal. Because these games are challenging to solve and usually display multiplicity of equilibria, most of the forward-looking part is restricted to the case with a very concentrated market which is composed by two representative consumers. Nevertheless, we find promising evidence that a large number of consumers facilitate tacit collusion. Indeed, to attract all the market, the deviating firm has to compensate consumers for the earnings they would have had by free-riding, i.e. let everyone switch but themselves. This decreases the deviation price and makes collusion easier. The main results of the paper are represented in Figure 3.1.2 where we sort the different types of consumers and establish when a collusive outcome is more likely to materialize.

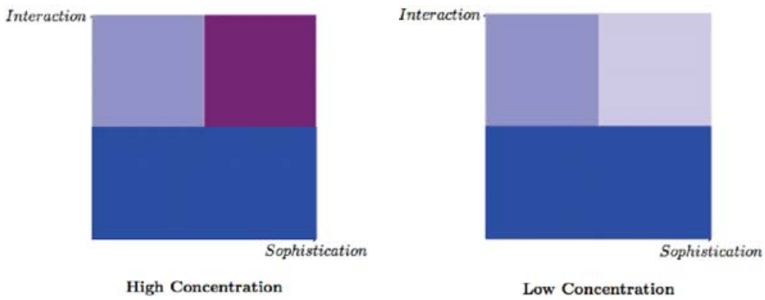

Figure 3.1.2: Sustainability of collusion with different type of consumers — Darker colors represents that full collusion is harder to implement.

There exists large empirical evidence that consumers do not always respond to price cuts the way predicted by classical models of Bertrand competition. Our paper contributes to the strand of the literature which explains this phenomenon by assuming the existence of switching costs associated to a change of supplier. Switching costs are the real or perceived costs incurred when changing supplier but which are not incurred by remaining with the current supplier.[3] Such costs are present in many sectors of the economy. For instance, in the IT sector they arise due to the use of proprietary formats and the

here.

[3]The definition is provided by the Office of Fair Trade.

resulting lack of compatibility across platforms, or the effort needed to learn how to use a new technology. Examples abound also in other sectors such as telecommunications and banking, but we can generalize the idea of switching costs as the savings generated by synergies resulting from a long-term relationship between a supplier and a retailer.

Whichever their type, switching costs create frictions on consumers' purchasing decisions due to the extra cost associated to a change of suppliers. In the present paper, we investigate the effect of such purchasing friction on the ability of firms to implement a tacit agreement. More precisely, on the effect that switching costs have on the likelihood to establish collusive prices. Economic theory has already established that the effect of switching costs on collusion is ambiguous. While switching costs reduce the incentives to deviate from collusion, by limiting the consumers' reactions to price cuts, it also reduces the severity of punishment. In the present paper, we show that the domination of one effect over the other depends crucially on the the type of consumers existing in the economy.

	Local calls			National long-distance calls			Calls to the United States		
	2000	2005	2010	2000	2005	2010	2000	2005	2010
EU-27	:	0.35	0.41	1.31	0.76	0.72	:	2.10	1.71
Belgium	0.49	0.57	0.63	1.74	0.57	0.63	5.95	1.98	2.17
Bulgaria	0.06	0.16	0.16	1.41	0.68	0.50	11.29	1.84	0.92
Czech Republic	0.55	0.64	0.65	1.66	1.30	0.65	:	2.36	2.34
Denmark	0.41	0.37	0.13	0.54	0.37	0.13	4.72	2.38	2.72
Germany	0.43	0.39	0.29	1.24	0.49	0.29	2.45	1.23	0.29
Estonia	0.14	0.25	0.25	0.71	0.25	0.25	10.26	2.56	2.31
Ireland	0.51	0.49	0.58	0.94	0.82	0.92	2.92	1.90	1.96
Greece	0.31	0.31	0.32	1.40	0.74	0.76	3.26	2.93	3.03
Spain	0.28	0.28	0.30	1.85	0.84	0.97	4.25	1.53	1.66
France	0.42	0.33	0.36	1.19	0.83	0.77	2.97	2.27	1.96
Italy	0.25	0.22	0.22	1.72	1.15	1.15	2.79	2.12	2.12
Cyprus	0.08	0.21	0.18	0.61	0.21	0.18	3.73	0.65	0.65
Latvia	0.35	0.35	0.36	1.01	1.01	1.04	5.81	5.83	5.98
Lithuania	0.26	0.39	0.39	1.07	0.79	0.79	11.96	4.06	4.12
Luxembourg	0.37	0.31	0.31	-	-	-	2.06	1.37	1.37
Hungary	0.34	0.41	0.46	1.22	1.07	1.12	4.24	2.93	2.40
Malta	:	0.25	0.25	-	-	-	:	1.76	1.90
Netherlands	0.30	0.33	0.60	0.42	0.49	0.60	0.78	0.85	1.82
Austria	0.69	0.49	0.54	2.30	0.59	0.54	4.32	1.90	2.20
Poland	0.36	0.36	0.51	1.49	1.30	1.02	10.72	3.79	1.24
Portugal	0.23	0.37	0.37	1.28	0.65	0.38	3.68	3.11	3.11
Romania	0.23	0.35	0.24	1.37	0.64	0.24	6.91	2.62	1.19
Slovenia	0.17	0.26	0.29	0.17	0.26	0.29	:	1.40	1.40
Slovakia	0.40	0.75	0.75	1.93	1.54	1.15	11.18	3.79	0.75
Finland	0.22	0.24	0.34	0.87	0.94	1.01	5.68	4.90	4.80
Sweden	0.30	0.30	0.29	0.30	0.30	0.29	1.12	1.08	0.89
United Kingdom	0.47	0.36	0.77	0.95	0.36	0.77	2.83	1.67	2.19
Japan	0.41	0.35	0.33	3.07	1.43	1.37	6.14	6.14	5.86
United States	0.09	0.08	0.08	0.43	1.02	0.67	-	-	-

(1) The indicator gives the price in euro of a 10-minute call at 11 am on a weekday (including VAT) for respectively a local call (3 km), a national call (200 km) and an international call to the United States; prices refer to August for 2000 and 2005 and to September for 2010; normal tariffs without special rates are used.
Source: Eurostat (online data code: isoc_tc_prc), Teligen

Figure 3.1.3: Retailing prices of different calling services in Europe.

As an illustration, consider the retailing prices of different calling services in Europe. The first two columns in Figure 3.1.3 show the prices of local and long-distance domestic calls for different years and different European countries. We observe two general price patterns for both services. For local calls, prices look homogeneous among countries and are rather low. For long distance calls however, we see that for some countries such as Germany, Sweden and Portugal, prices have severely decreased in the last decade to a level similar to local calls. In countries like Spain, France, Italy and Finland however, prices for long-distance calls are still substantially higher than local calls. Moreover, this pattern does not apply when we consider calls in the direction of the United States represented in the last column, where a substantial price reduction has occurred in almost all countries.

Our paper explains these differences by switching costs and potential tacit collusion. In general, the local call service is characterized by heavy usage while the one for long distance is of milder use. In our model we approximate usage with the level of interaction between consumers and firms, and we show that implementability of collusion is affected by the degree of interaction. High prices are therefore not surprising in the market for long-distance calls as these markets are the more susceptible to sustain collusion. We also consider the degree of demand concentration as a possible explanation for the implementability of implicit contracts.[4]

Due to the link between buyer's structure and the likelihood to sustain a collusive outcome, our paper has clear policy implications. A competition authority should take into account the type of consumers in the industry to examine the possibility of firms to implement implicit contracts. In general, we find that collusion is harder to sustain when consumers are more sophisticated. As a result, the aim of an antitrust authority should be to dedicate resources on educating consumers. In other words, consumers should be aware of how their current purchasing decisions affect future prices.

The rest of the paper is organized as follows. After briefly discussing the existing literature, section 3.3 presents the model. Section 3.4 deals with two forward-looking consumers, while section 3.5 treats the case of a continuum of consumers. Later, in section 3.6 we treat alternative specifications for the consumers, namely their sophistication, lifespan, and the possibility of increasing switching costs. Finally, we examine the case of increasing loyalty in section 3.7.

3.2 Related Literature

Following much of the literature on switching costs, we consider that only short term contracts are available and we do not allow firms to price discriminate between consumers. As a result, at each period,

[4]Unfortunately, our model will not explain market structure as it only tries to predict market prices with respect to the type of consumers in the market. Indeed, for most of the cases a price lower than full collusion appears whenever the market structure is asymmetric and one of the firms captures the whole market.

each firm faces the trade-off between setting a low price to attract the consumers of the rival or a high price to extract rents from the captive consumer.[5] This trade-off is present in many two-period models of switching costs such as Farrell [21], Klemperer [31], Klemperer [30], Klemperer [29], and Padilla [43]. A two-period setting however is inadequate to study the firms' incentives to tacitly collude. In this regard our model is closer to models with an infinite number of periods like Farrell and Shapiro [20], Beggs and Klemperer [5], Rosenthal [47], Rosenthal [48], Padilla [44], and Anderson et al. [2].[6] The vast majority of those papers study the implications that switching costs have on competition and center attention on the steady state equilibrium price and market shares. The last two works study the likelihood of implementing collusive prices by firms when consumers have switching costs. However, our analysis differs from theirs in the following aspects: i) our work considers both short-run and long-run consumers; ii) instead of working with an overlapping-generation model, our long-run consumers live infinitely; and iii) we draw implications regarding the level of sophistication of our long-run consumers.[7] Therefore, our analysis is centred on the type of consumers existing in the industry and more specifically on the live-span and the level of sophistication of those consumers. We also consider how the concentration of consumers in the industry affects the implementability of collusion and how our results change if we allow the switching costs of consumers to increase with captivity, i.e. to become increasingly loyal.

Assuming infinite interactions between firms and consumers allow us to examine a rich set of consumer strategies. In our model, consumers decisions have a large impact on the likelihood of collusion. By considering that all players in our game are strategic and forward-looking, our work relates to the analysis of Biglaiser and Crémer [7] and Biglaiser et al. [8]. However, those papers investigate different issues than ours. The first is the application of the Folk theorem when consumers have switching costs, and the second studies the effects that heterogeneous switching costs have on competition. None of this two previous works study the case of sustainability of a collusive outcome. Finally, our analysis of consumers as short-run players also relates our paper to the existing literature on mixed pricing in oligopoly, as in Shilony [49] and Narasimhan [39]. A mixed strategy equilibrium appears due to the fact that the profit function of firms fail to be continuous. We applied Shilony's methodology to recover our results but we obtained an additional equilibrium in mixed strategies if indifferent consumers choose to switch. Hence, we obtain multiplicity of equilibria on some range of switching costs by not assuming any specific tie-breaking rule for consumers.

[5] A recent study, Cabral [10], gives an analysis of the case in which firms can price discriminate between captive and non-captive consumers. Similar examples are found in Caminal and Matutes [12], Eber [17] and Nilssen [40].

[6] Other dynamic models with switching costs are Fabra and García [19], Lewis and Yildirim [33], To [53], Rhodes [46], and Villas-Boas [54].

[7] The previous papers have not been able to deal with the case when consumers are forward-looking. Anderson et al. [2] highlights a mistake in Padilla [44] studying this case.

3.3 Model

Two firms A and B produce a non durable homogeneous good at zero cost. On the demand side, two consumers $j = 1, 2$ wish to buy a single unit of the good at each period and obtain a gross utility of 1. Later in section (3.5) we consider the case of a continuum of consumers. Consumers are captive with degree $S \in [0, 1]$ to firm A or B. The degree of captivity in our model is exogenous and common to each consumer. It represents the magnitude of switching costs. If a consumer captive to one firm buys from the other she incurs an extra cost of S. We denote by C_j^t the identity of the firm to which consumer j is captive at date t and we assume that at the beginning of the game, firms equally share the market. The state of the world ξ_i^t represents the number of consumers captive to firm i at the beginning of period t.

In our model, firms interact for an infinite number of periods and each period is composed of two stages. In stage 1, there is a Bertrand-type competition where firms set their price simultaneously and non cooperatively $p_i^t \in \mathbb{R}$ for $i = A, B$. In stage 2, consumers make their purchasing decisions $d_j^t = \{A, B, \emptyset\}$, where the empty set ensures that purchasing is voluntary. Finally, our game is one of complete information since firms and consumers are able to observe the full history of prices and purchasing decisions.

Consumers' behaviour : Consumer behaviour depends on the type of consumers in the economy. The first aspect concerns their lifespan or level of interaction with the firms. Here, we consider two extreme cases, consumers are either short-run or long-run depending on whether they can purchase only once or infinitely. We further differentiate consumers by their sophistication. Consumers can be either short-sighed or forward-looking. When consumers are short-sighted, they maximize current utility and buy from the firm offering the lowest price once discounted for switching costs. When consumers are forward-looking, they anticipate that their purchasing decisions affect future equilibrium prices. Accordingly, they behave strategically and take current purchasing decisions anticipating future prices and switching costs.

Moreover, the analytical difference between short-run and long-run consumers lies on the way captivity changes over periods. With short-run consumers, captivity remains unchanged regardless of past actions: old consumers leave the market at the end of each period and are replaced by new ones in the next period. With long-run consumers, the captivity of a consumer changes according to her previous purchasing decision: $C_j^t = d_j^{t-1}$. If a consumer does not effectuate any purchase at time t, i.e, $d_j^t = \emptyset$, then her captivity remains unchanged $C_j^t = C_j^{t-1}$. Table 3.3.1 summarizes the type of consumers considered and its strategic effect in the model.

Table 3.3.1: Type of consumers

Live-span	Sophistication	Decision	Captivity
Short-run	-	Current utility	Fixed
Long-run	Short-sighted	Current utility	Variant
Long-run	Forward-looking	Discounted utility	Variant

Firms' behavior : At each period, firms compete simultaneously and non cooperatively in prices. We assume that firms do not observe the identity of consumers who are captive, and so price discrimination is not feasible.[8] We also consider that firms are not capacity constrained and for any market price, they always fulfil demand. Therefore, at the beginning of each period, firms simultaneously post prices to maximize their discounted expected profit:

$$\Pi_i = \sum_{t=1}^{\infty} \delta^{t-1} \pi_i^t(p^t, d^t), \quad i = A, B, \tag{3.3.1}$$

where p^t is the firms' action profile, d^t the consumers' choice at time t and δ the discount factor.

Simple strategy profiles : To obtain the equilibrium sequence of prices of the game, we restrict attention to simple strategy profiles. In our game, a simple strategy profile consists of an initial prescribed outcome $\mathbf{p}(c)$ and a punishment outcome $\mathbf{p}(i, \xi_i)$ for each firm $i = A, B$ and each state of the world $\xi = 0, 1, 2$, where $\mathbf{p}$ is an outcome path $(p^1, p^2, \ldots)$ that consist of a vector of two prices for each period. We focus on symmetric stationary equilibria, hence:

- $\mathbf{p}(A, \xi) = \mathbf{p}(B, \xi)$ for all ξ and we can omit the firm index.
- If firm i is in state ξ_i, firm $-i$ is necessarily in state $\xi_{-i} = 2 - \xi_i$, hence plays $\mathbf{p}(2 - \xi_i)$.
- In any path $\mathbf{p}(\xi)$, the firm indefinitely posts a price $\mathbf{p}(\xi)$.

As a result, a simple strategy profile in our duopolistic game is fully characterized by $(\mathbf{p}(c), \mathbf{p}(\xi)_{\xi=0,1,2})$. The game starts with the initial path $\mathbf{p}(c)$ and players continue with this path until a deviation takes place, where a deviation might come from either firms or consumers. We define the initial path as the collusive one. If a deviation by any firm i has occurred resulting in the state of the world ξ, the game switches to the

[8] This assumption might seem quite restrictive because nowadays firms can keep track of the consumers that have previously purchased. Indeed, some literature studies behaviour price discrimination where firms set different prices to captive and prospective consumers. Some justification of the impossibility to price discriminate might be a regulation banning such a practice, or firms not retailing directly the product, e.g. google android vs. iphone.

path in which firm i plays $\mathbf{p}(\xi)$ and firm $-i$ plays $\mathbf{p}(2-\xi)$ forever. Therefore restrict attention to collusive agreements supported by grim-trigger strategies. We have that:

Definition 5. *Full collusion is implemented at equilibrium by a simple punishment path* $\mathbf{p}$ *if* $\mathbf{p}(c) = 1$ *and* $\mathbf{p}(\xi)_{\xi=0,1,2}$ *constitute a stationary equilibrium.*

By stationary equilibrium, we mean that in each possible state of the world, either in the collusive path or if firm i plays $\mathbf{p}(i, \xi_i)$ and firm $-i$ plays $\mathbf{p}(-i, 2-\xi_i)$, then no firm has any incentive to deviate, but also no consumer wants to switch. Therefore, our equilibrium concept is stationary sub-game perfect equilibrium (SSPE). Finally, we consider:

Assumption 4. ***(Tie-breaking rule):*** *Any firm deviates from a stationary path if it gets strictly higher profits.*

We do not assume any specific tie-breaking rule coming from the consumer. Indeed, for short-run consumers, we find different equilibria depending on whether a consumer decides to switch or not when she is indifferent.

3.4 Long-run Consumers

With long-run consumers, captivity changes and it is equal to the firm from which a consumer has last purchased, i.e. $C_j^t = d_j^{t-1}$ if $d_j^{t-1} = A, B$.[9] The discounted utility that consumer j obtains by purchasing from any firm i at time t is:

$$U_j(\cdot) = \sum_{t=1}^{\infty} \delta^{t-1} \left[1 - p_i^t - 1(d_j^t \neq C_j^t)S\right]$$

and consumers pay switching costs as long as their purchasing decision is different from the one made in the last period. As we want to characterize when a full collusive outcome can be implemented, in what follows, we obtain the punishment path $\mathbf{p}(\xi)$ coming from any deviation that moves the game to any state $\xi = 0, 2$. We consider the case where the punishment path $\mathbf{p}(\xi = 1)$ is also full collusion.[10]

3.4.1 The Incumbency Game

Here, we study the equilibrium of the continuation game when one of the firms sold to both consumers in the last period. This situation may arise either because one of the firms have deviated from the collusive outcome and attracted the consumer of the rival, or one of the consumers has switched without any deviation by a firm. Therefore, we want to characterize the path $\mathbf{p}(\xi)$ for $\xi = 0, 2$.

[9]If $d_j^{t-1} = \emptyset$, we find the largest τ with $d_j^\tau \neq 0$ and $C_j^t = d_j^\tau$.

[10]Whether other equilibria exist with a price lower than $\mathbf{p}(\xi = 1)$ is a question that we leave for further research.

The sub-game when one of the firms has both consumers is similar to a market structure with an incumbent and a potential entrant, in the sense that if any consumer decides to buy the product offered by the entrant, he has to incur a switching cost. Let us call I the firm with two captive consumers, and E the firm with none. Therefore, the decision sub-game played by the consumers captive to I and facing prices (p^I, p^E) is represented in the Table 3.4.1. Each cell contains the present discounted payoffs originating

Table 3.4.1: Payoffs of the Incumbency game

	St	Sw
St	$\left(\frac{1-p^I}{1-\delta}, \frac{1-p^I}{1-\delta}\right)$	$(1-p^I, 1-p^E-S)$
Sw	$(1-p^E-S, 1-p^I)$	$\left(1-p^E-S+\delta\frac{1-p^I}{1-\delta}, 1-p^E-S+\delta\frac{1-p^I}{1-\delta}\right)$

from each action. Sw and St stand for switching and not switching.[11]

The following lemma states the price set by the incumbent. The equilibrium prices depend on the type of coordination game played by consumers and their beliefs about the action of the other buyer.[12]

Lemma 1. *Restricting attention to our simple strategy profile, any incumbent price lies in the interval*

$$p^I \in \left[\max\left\{0, S - \frac{\delta}{1-\delta}\right\}, \frac{\delta + S(1-\delta)}{1+\delta}\right].$$

There is a continuum of prices that the incumbent may set. This multiplicity comes from the coordination game played by consumers and depends on whether consumers coordinate their actions and on their beliefs about the actions of the other consumer. We detail the proof of Lemma 1 in Appendix C.1.

Consumers coordinating on the Pareto dominant equilibrium seems a reasonable assumption. For instance, advertising by the entrant could provide a focal point to consumers and entice them both to switch. Failure to coordinate however is very hurtful to consumers and one can argue that consumers switch only if doing so is an equilibrium in dominant strategies. If so, the equilibrium incumbent and entrant prices change. To attract consumers, the entrant needs to post a lower price and the incumbent is able to set a higher price to keep her consumers. There also exists an equilibrium in which both consumers switch and this gives them an utility lower than if they coordinate on sticking.

[11] Since we focus on fully collusive equilibria and because we impose no particular punishment when $\xi = 1$, we have $p(c) = p(1) = 1$. Note that if both consumers switch to the entrant, they anticipate that in the continuation game they will pay the price of the incumbent p^I.

[12] We find different equilibrium prices than in Biglaiser and Crémer [7]. In particular, we focus on an equilibrium in which the potential entrant is able to set negative prices. Equilibria in Biglaiser and Crémer [7] rely on more complex punishment structures than the one in our model.

3.4.2 Collusive Equilibrium

We now consider the collusive path and characterize the conditions under which full collusion is implementable. To maintain full collusion, the following incentive constraints need to be satisfied:

- No firm wants to deviate from collusion.

$$\frac{1}{1-\delta} \geq 2\left(\tilde{p} + \delta \frac{p(2)}{1-\delta}\right), \qquad (IC_f)$$

  where $\tilde{p}$ is the optimal deviation price.

- No consumer wants to unilaterally deviate from the collusive path.

$$U_j\left(\mathbf{p}(c)\right) \geq 1 - p(c) - S + \delta \times U_j\left(\mathbf{p}(2)\right), \qquad (IC_c)$$

- A firm deviating from the collusive path charges the highest possible price $\tilde{p}$ which attracts the consumer of the competitor (IC_a).
- After any deviation, one of the incumbency game equilibria characterized in Lemma 1 is played.

Note that (IC_a) tells us that by setting a deviation price p^d, the consumer of the other firm must be enticed to switch to the deviating firm:

$$1 - p^d - S + \delta \frac{1-p(2)}{1-\delta} \geq 0, \qquad (IC_a)$$

where $\tilde{p}$ corresponds to the deviation price such that (IC_a) binds. It is easy to check that the consumer of the deviating firm does not want to switch. Substituting $\tilde{p}$ in (IC_f) gives us a necessary condition for collusion and ensures that no firm unilaterally deviates from the collusive outcome:

$$\frac{1}{1-\delta} \geq 2\left(1 - S + \delta \frac{1-p(2)}{1-\delta} + \delta \frac{p(2)}{1-\delta}\right) \Leftrightarrow \delta \leq 1 - \frac{1}{2S} = \bar{\delta}(S).$$

Note that the incumbency price $p(2)$ does not have any effect on this condition. Therefore, the price paid in the punishment phase does not affect the firms' incentives to sustain a collusive outcome. Because consumers are forward-looking, they accept to switch if the deviating firm reimburses the switching cost minus what they save by buying from an incumbent forever. The deviation price then takes into account future rents enjoyed by consumers and depends on the equilibrium played in the incumbency game. As a result, the incumbency price disappears from the firms' incentive compatibility.

Since forward-looking consumers behave strategically, it is left to verify that no consumer wants to deviate from full collusion. The incentive constraint of consumers (IC_c) is harder to fulfil when the incumbency price decreases. Hence, non deviation of consumers from collusion occurs whenever the incumbency prices is the lowest possible. Using the range of possible prices given by Lemma 1, we get:

$$0 \geq -S + \delta \frac{1 - \underline{p}(2)}{1 - \delta} \Leftrightarrow S \geq \delta \frac{1 - \left(S - \frac{\delta}{1-\delta}\right)}{1 - \delta} \Leftrightarrow S \geq \frac{\delta}{1 - \delta}.$$

It suffices to check if this condition is compatible with a value of δ ensuring that no firm deviates, which is always the case as long as $\delta \leq 1/2$. Therefore, the above necessary condition to sustain a full collusive outcome is also sufficient. Since for this range of the parameters, no firms or consumers deviate from full collusion, the punishment path $p(1)$ constitutes also a stationary equilibrium. With all these results, we get the following proposition.

Proposition 8. *When consumers are forward-looking, full collusion is sustainable if and only if* $\delta \leq \max\left\{0, 1 - \frac{1}{2S}\right\}$ *and it is independent of the punishment path played after a deviation occurs.*

Figure 3.4.1 displays the set of parameters sustaining full collusion:

$$C_{FL} : \left\{(\delta, S) \mid \delta \leq \max\left\{0, 1 - \frac{1}{2S}\right\}\right\}$$

We see that full collusion is only sustainable for large value of switching costs together with a low discount factor: the more players value future outcomes, the larger is the deviation price and the bigger the rents from deviating. The opposite occurs whenever the discount factor is small and switching costs are large.

Remark 1. *The equilibrium played in the incumbency game is only used to verify that consumers do not have incentives to break the collusive outcome and the likelihood to maintain collusion is neutral on the equilibrium price arising after deviation. This result comes from our assumption of identical discount factor for firms and consumers.*

We now consider a general case in which the discount factors of firms and consumers do not coincide. We see that both discount factors have opposing effects on sustaining a collusive outcome and the equilibrium played in the incumbency game influences the sustainability of collusion.

Heterogeneous Discount Factor

Again, in order to implement a collusive outcome, we need to verify that no firm or customer wants to unilaterally deviate from collusion. However, now the discount factor of firms and consumers may differ

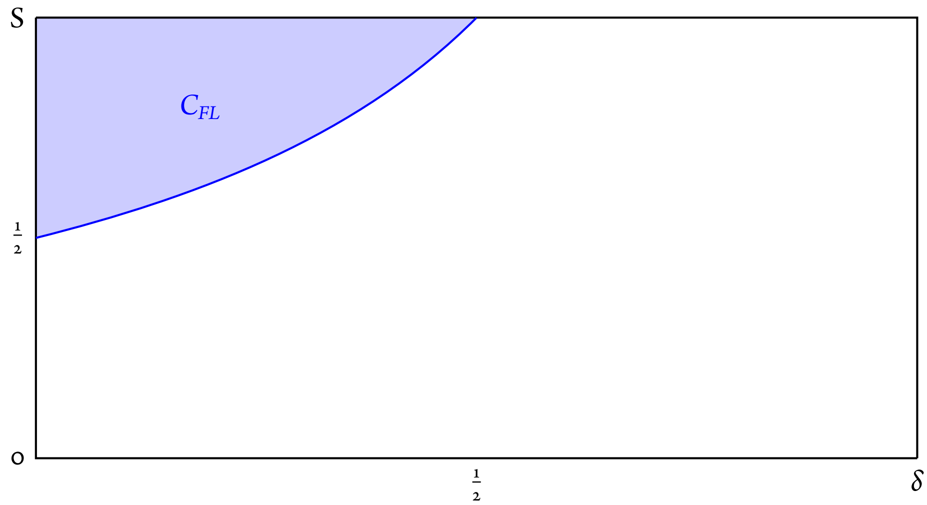

Figure 3.4.1: Full collusion in the case of two forward-looking consumers — C_{FL} is filled in blue.

$\delta_f \neq \delta_{co}$. The price in the incumbency game coincide with the one presented in Lemma 1 as this is independent on the weight that consumers put on the future.[13] By solving (IC_a) as in the previous section, we obtain that the maximum deviation price such that the consumer of the rival is attracted equals:

$$\tilde{p} = 1 - S + \delta_{co} \times \left(\frac{1 - p(2)}{1 - \delta_{co}} \right),$$

Introducing this into the incentive constrain of firms, we get the necessary condition to implement a collusive outcome:

$$\frac{1}{2(1-\delta_f)} \geq 1 - S + \frac{\delta_{co}}{1-\delta_{co}} + p(2) \times \left(\frac{\delta_f - \delta_{co}}{(1-\delta_f)(1-\delta_{co})} \right). \qquad (IC_f)$$

With different discount factors, the equilibrium played in the incumbency game has an effect on the necessary condition to implement collusion and its intensity and sign is measured by the term $\mathcal{D} = \frac{\delta_f - \delta_{co}}{(1-\delta_f)(1-\delta_{co})}$. If the consumers discount factor is higher than the firms discount factor, a larger incumbency price makes collusion harder to sustain. The opposite happens when firms have a greater valuation for the future than consumers. The intuition comes from the analysis of the optimal deviation price. As long as consumers have a larger valuation for the future, firms are able to set a large deviation price. However, the magnitude of this price decreases with the equilibrium incumbency price posted after

[13] The price in the incumbency game depends only on the discount factor of firms $p(2) \in \left[\max\left\{0, S - \frac{\delta_f}{1-\delta_f}\right\}, \frac{\delta_f + S(1-\delta_f)}{1+\delta_f}\right]$.

deviation. The rents firms can get at the moment of deviation are low whenever the incumbency prices are large. However, when the valuation for the future is larger for firms than for consumers, a large incumbency price has a low effect on the deviation price and an important effect on the continuation payoffs after deviation. Collusion is harder to sustain with an increase of the incumbency price.

The firms' incentive constraint needs to be completed by the fact that no consumers wants to deviate from the collusive path.

$$S \geq \delta_{co} \times \left(\frac{1 - p(2)}{1 - \delta_{co}} \right).$$

To obtain a solution of the game, we consider that the equilibrium of the incumbency game is the one arising from perfect coordination between consumers $p(2) = (1 - \delta_f)S$.[14] Hence, the optimal deviation price is $\tilde{p}\left(\delta_{co}, \delta_f, S\right) = \frac{1-(1-\delta_{co}\delta_f)S}{1-\delta_{co}}$. Solving for the incentives constraint of firms and consumers, full collusion constitutes an equilibrium if the following two conditions are fulfilled:

$$\delta_f \geq \frac{-1 + 2S + \sqrt{1 - 2(1 - \delta_{co})S}}{2S} = \underline{\delta}_f(\delta_{co}, S) \quad \text{and} \quad \delta_{co} < \frac{S}{1 + \delta_f S} = \bar{\delta}_{co}(\delta_f, S).$$

In Figure 3.4.2, we draw the region of parameters sustaining a full collusive outcome. When the discount factor of consumers increases, the collusive region is reduced since the deviation price is larger. Because consumers obtain rents from switching, the deviation prices cashes in this effect. We remark that the region where full collusion is sustainable is not monotone for a positive discount factors of consumers. This non-monotonicity is due to the sign of the term $\mathcal{D}$. We then obtain that, for a given discount factor of consumers, the region of full collusion is decreasing with δ_f when $\delta_{co} \geq \delta_f$ and increasing with δ_f when $\delta_{co} < \delta_f$. Moreover, we see that this pattern is reversed again for a $0 < \bar{\delta}_{co} < \delta_f$. In this case the incentive constraint of consumers is binding and for some values of switching costs, consumers decide to break the collusive outcome.

[14]Notice that there exist as many equilibria in the collusive game as in the incumbency game — i.e. a continuum — since the incentive condition of firms depends on the equilibrium price of the incumbency game.

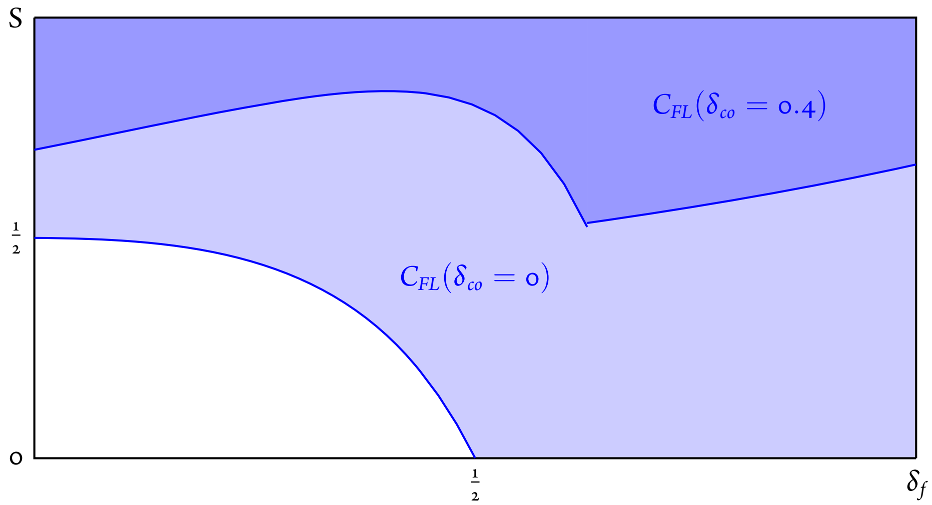

Figure 3.4.2: Full collusion with two forward-looking consumers and different discount factors for firms and consumers — The blue area represents the region where full collusion is sustained for a given value of consumer's discount.

3.5 Continuum of Consumers

We now generalize the model to a continuum of consumers and look for conditions supporting a full collusive equilibrium. We find that whenever consumers are forward-looking, collusion is much easier to sustain in the case of a continuum of consumers than when there are two consumers in the market. Because consumers do not have mass, their individual purchasing decisions do not have any effect on equilibrium prices and they can free-ride on the decisions of their peers. In order to ensure switching, a deviating firm has to pay the consumer for this possible gains and this reduces drastically the profits from deviation.

3.5.1 Description of the Game

The structure and timing is the same as in the case of two consumers but now there is a continuum of consumers of mass 2 in the market. As before, each firm is endowed with half of the market and we focus on stationary symmetric equilibria. Our simple strategy profile is summarized by an initial collusive outcome $\mathbf{p}(c)$ and a punishment outcome $\mathbf{p}(\eta) : \eta \in [0, 2] \to [-\infty, 1]$. The state of the world is now a continuous variable η which represents the mass of consumers detained by a firm at the beginning of the period.

As before, we restrict attention to grim-trigger strategies and each firm initially posts a collusive price

$\mathbf{p}(c)$. In the case of a deviation resulting in firm i serving a mass η of consumers, the firms move to the stationary punishment phase in which firm i indefinitely posts $\mathbf{p}(\eta)$ and firm $-i$ posts $\mathbf{p}(2-\eta)$. The following assumption is necessary for the resolution of the model.

Assumption 5. ***(Continuity):*** *$p(\cdot)$ is continuous over $[0, 2]$.*

We justify this assumption by the fact that firms cannot monitor infinitesimal deviations from consumers. Therefore their price functions must be continuous. This also comes from the fact that consumers do not have "strategic mass" in the sense that their individual purchasing decisions do not have any effect on the equilibrium prices. We finish the presentation of the model with a conjecture and a definition.

Conjecture 1. ***(Restriction over the strategy sets):*** *When deviating, a firm aims at making all consumers switch.*

Definition 6. *Full collusion is implemented in equilibrium by a simple punishment path* $\mathbf{p}$ *if* $p(c) = 1$ *and* $p(\eta) : \eta \in [0, 2] \to [-\infty, 1]$ *constitute a stationary equilibrium.*

With a purpose of tractability, we look at equilibria in which $p(c) = p(1) = 1$, i.e. if the game moves back to a situation in which firms equally share the market, full collusion takes place.[15] We proceed as in the case of two representative consumers: first we obtain the equilibrium of the incumbency game and then we examine when full collusion can be implemented. In this specification of the model, we consider that the discount factors of the firms and of the consumers are the same.

3.5.2 The Incumbency Game

Let p^{EI} denote the deviation price by the entrant such that consumers switch. p^{EI} is such that, for every consumer, the profits obtained from switching to the entrant are higher than the ones from staying with the incumbent. When there is a continuum of consumers, individual purchasing decisions do not have any effect on the market equilibrium price. As a result, as we assume that all consumers must switch, the incentive compatibility constraint for a single consumer is :

$$1 - p^{EI} - S + \delta\frac{1 - p(2)}{1 - \delta} \geq 1 - p(2) + \delta\frac{1 - p(0)}{1 - \delta}, \qquad (IC_a^{\infty}) \qquad (3.5.1)$$

where the right-hand side represents what a consumer gets by sticking to the incumbent provided that all the other consumers have switched. We call the last term the free-riding payoff. This is what a consumer

[15] This is in full contradiction with Fabra and García [19]. In the second Lemma of their paper, they show that the market price decreases with market symmetry. We, however, look at the likelihood to implement implicit contracts. Hence a more symmetric market implies a larger probability to maintain a collusive outcome.

obtains from not switching given that all other consumers have switched $U^{FR} = \delta \times \left(\frac{1-p(0)}{1-\delta}\right)$.[16] The entrant chooses the highest deviation price p^{EI} such that the previous expression is binding and ensures that all consumers switch.

$$\bar{p}^{EI} = p(2) - S - \frac{\delta}{1-\delta}(p(2) - p(0))$$

By substituting the previous expression to the constraint on negative profits for the deviating entrant, we obtain the upper-bound on the incumbency price $p(2)$:

$$\bar{p}^{EI} + \frac{\delta}{1-\delta}p(2) \leq 0 \Leftrightarrow p(2) \leq S - \frac{\delta}{1-\delta}p(0). \tag{3.5.2}$$

Finally, as an equilibrium in our game is such that no consumer switches from any path, the constraint on no unilateral deviation from consumers needs to be satisfied and this gives another upper bound for p(2)

$$\frac{1-p(2)}{1-\delta} \geq 1 - p(0) - S + \delta\frac{1-p(0)}{1-\delta} \Leftrightarrow p(2) \leq p(0) + S(1-\delta). \tag{IC_u^∞}$$

Condition (3.5.2) is the incentive constraint saying that the entrant and incumbent have no profitable deviation. Condition (IC_u^∞) represents the incentive constraint saying that no consumer unilaterally deviates. This represents an extra constraint with a continuum of consumers since it is always slack with two representative consumers. With a continuum of consumers, the incumbency price is determined by the lower-envelope of the two previous conditions. Moreover, both conditions, need to be binding. Consider a situation where constraint (IC_u^∞) is not binding, then there is an upward profitable deviation for the incumbent. Suppose both firms post prices such that (3.5.2) holds with equality, then $p(2)$ is not a best-response to $p(0)$. By raising $p(2)$, the incumbent does not lose any consumers since (IC_u^∞) is not binding. Condition (3.5.2) merely states that stable prices in the incumbency game are impervious to deviation by the entrant. We therefore need (3.5.2) to hold, but we also need (IC_u^∞) to hold with equality. All this leads to the following lemma whose proof is direct:

Lemma 2. *The pair of prices $(p(0), p(2))$ such that:*

$$p(2) = \begin{cases} S(1-\delta) + p(0) & \text{if} \quad p(0) \leq \delta(1-\delta)S \\ S - \frac{\delta}{1-\delta}p(0) & \text{if} \quad p(0) \in \left(\delta, \frac{1}{2-\delta}\right] \times (1-\delta)S \end{cases}$$

constitutes a stationary equilibrium of the incumbency game.

Compared to the case of two representative consumers, the only relevant constraint refers to the

[16]Notice also that if this condition is satisfied, then coordinating on switching yields higher utility than coordinating on sticking. In the duopsony case, a consumer obtains zero profits if he does not switch.

behaviour of the entrant and the form of the optimal deviation price completely determines the maximum incumbent price. Now, the optimal deviation price is a function of both the incumbent and the entrant prices. Moreover, considering our continuity assumption, taking $p(0) < 0$ would violate the participation constraint of the entrant. Indeed, by continuity, there exists a mass $\eta > 0$ of consumers such that $p(\eta) < 0$ and the firm would make negative profits. Hence, to avoid this, the lowest value that $p(0)$ can take is zero, and the incumbent price in this case is $p(2) = (1 - \delta)S$.

Remark 2. *Our result relies heavily on the assumption that the entrant tries to attract all consumers or none. However, it is hard to think about possible deviations that attract only part of the market because a tremendous coordination on the side of consumers might be required. Also, the continuity assumption plays an important role here since it helps us to easily model what happens in case of an unilateral deviation. In contrast, intermediate models with N consumers are harder to solve as we need to explicitly specify the pricing strategies in all states of the world.*

3.5.3 Collusive Equilibrium

Suppose that a firm wants to break the collusive outcome. To do so, the deviating firm posts the highest price to make all consumers switch. Therefore condition (IC_a^∞) needs to be fulfilled:

$$1 - p^d - S + \delta \frac{1 - p(2)}{1 - \delta} \geq 1 - p(c) + \delta \frac{1 - p(0)}{1 - \delta}.$$

Hence, the optimal deviation price is the one making the previous expression binding and is given by $\tilde{p}^\infty = 1 - S - \frac{\delta}{1-\delta}(p(2) - p(0))$. Firms do not deviate from full collusion if:

$$\frac{1}{2}\frac{1}{1-\delta} \geq 1 - S - \frac{\delta}{1-\delta}(p(2) - p(0)) + \frac{\delta}{1-\delta}p(2) \iff \frac{1}{2}\frac{1}{1-\delta} \geq 1 - S + \frac{\delta}{1-\delta}p(0). \qquad (IC_f^\infty)$$

As in the duopsony model, the firms incentive to sustain full collusion does not depends on the equilibrium price $p(2)$ played on the incumbency game. This price is incorporated in the optimal deviation price. Moreover, it is easy to check that unilateral switching from consumers is never profitable, i.e. in the coordination sub-game, sticking to the collusive path is always an equilibrium for consumers.[17] The constraint (IC_u) is therefore always fulfilled. Moreover, the condition for full collusion only depends on the primitives and $p(0)$. The higher $p(0)$, the harder it is to sustain collusion because the optimal deviating price increases with $p(0)$. Hence, without further restrictions, any price $p(0)$ supporting an incumbency game equilibrium defines a set of parameters under which full collusion is implementable. It

[17]This is given by the continuity assumption on prices. A consumer that unilaterally deviates pays the same collusive price plus the cost of switching.

is now natural to consider the worst punishment phase, i.e. the punishment phase under which $p(0)$ is minimal. Using the results of Lemma 2, we choose $p(0) = 0$. The region where a full collusive outcome is implemented is introduced in the following proposition.

Proposition 9. *With a continuum of consumers, the largest set under which collusion is implementable is given by* $C_\infty = \left\{ (\delta, S) \mid \delta \geq \max\left\{ \frac{1-2S}{2(1-S)}, 0 \right\} \right\}$.

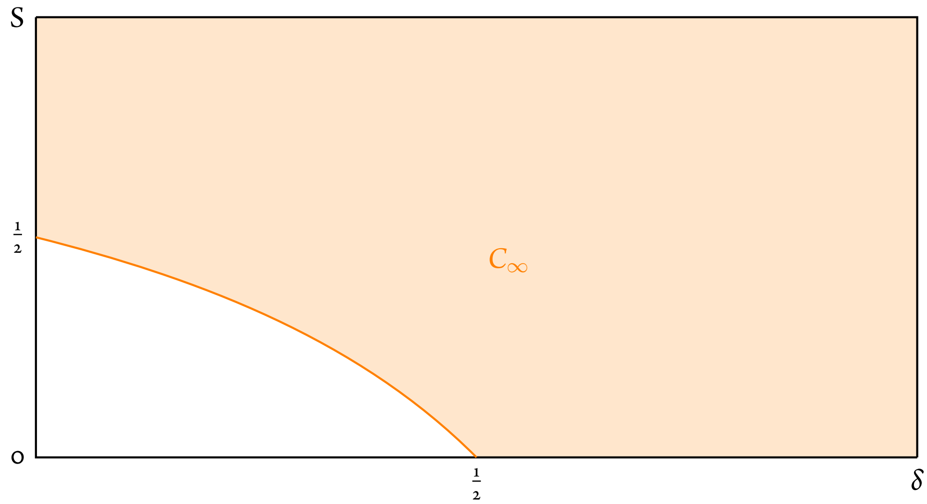

Figure 3.5.1: Full collusion in the case of a continuum of consumers — C_∞ is filled in orange.

Figure 3.5.1 depicts the region sustaining full collusion in this case. Clearly, the area where full collusion under a continuum of consumers is sustained is wider than in the case with two representative consumers. With a continuum of consumers both switching costs and the discount factors positively affect the implementability of implicit contracts. This is because now full collusion depends on the entrant price, which determines the free-rider profit. In the two consumers case, free-riding is equivalent to status quo and brings nothing to consumers. Here, free-riding implies playing the low entrant price, and getting large rents. To prevent this, the deviating firm must subsidize the consumers by the amount of this surplus, which makes collusion easier. Additionally, because with a continuum of consumers, each individual consumer does not have any strategic mass, the constraint representing unilateral deviations is always slack. In the following subsection, we consider that consumers might act cooperatively and coordinate.

Collusion with coordinating consumers

In the previous analysis, we did not consider the possibility for consumers to coordinate on the Pareto dominant equilibrium. One of the reasons we obtained a large region with possible collusion is because we did not consider the possibility that consumers may act cooperatively to break the collusive outcome. This might be feasible in situations where consumer associations or even deviating firms provide focal points on equilibria.

Compared to the previous analysis, we have to add an additional constraint coming from coordinating consumers which is similar to the one introduced in section 3.4.2 where we consider only two consumers. Coordinating consumers decide not to break full collusion if:

$$0 \geq -S + \delta \times \left(\frac{1 - p(2)}{1 - \delta} \right)$$

Notice that due to the relation between the entrant and incumbent prices shown in Lemma 2, setting the punishment path $p(0) = 0$ might not be the best course of action as it gives the consumers a great incentive to coordinate on breaking the collusive outcome. Indeed, in this case, the incumbency price is low. By raising $p(0)$, we relax the consumers' incentive constraint and can get a larger collusive region. Lemma 2 gives us the maximal $p(2)$ which can be implemented in equilibrium. This price is equal to $p(2) = (1 - \delta^2)S$. Figure 3.5.2 represents the collusive regions when the punishment path is set at the two extreme values of $p(2)$. We see that neither one nor the other solutions unambiguously increase the implementability of collusion. Notice however, that both solutions allow to implement collusion on a broader set of parameters than in the case of two representative consumers as the deviation price with a continuum of consumers is always smaller.

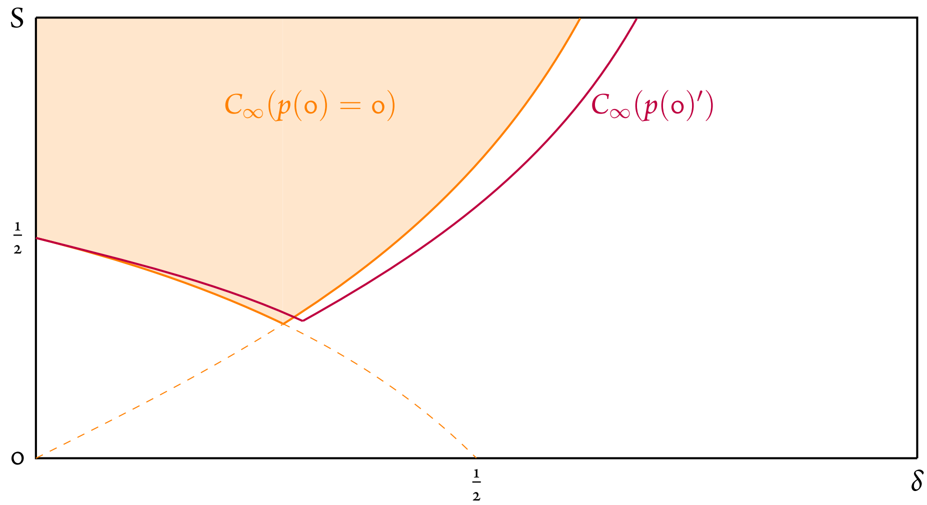

Figure 3.5.2: Full collusion in the case of a continuum of consumers when they coordinate on the Pareto dominant equilibrium — There is an additional constraint stemming from the ability to coordinate. The resulting collusive set is filled in orange. When $p(0)' = \delta(1 - \delta)S$, the higher incumbent price relax this constraint a little but gives more bite to the firms' constraint. The lower boundary of the collusive set is represented by the purple curve.

3.6 Variations in Consumers' Sophistication and Lifespan

3.6.1 Short-sighted consumers

When consumers are short-sighted, they do not anticipate that current purchasing decisions will have an effect on future equilibrium prices. Hence, they base their purchasing decisions on maximizing current utility. Because consumers do not behave strategically, we do not have to consider their beliefs on the decisions of the other consumers. Hence, the unique equilibrium in the incumbency game is identical to the one with full coordination $p(2) = (1 - \delta)\, S$.

The conditions to sustain a collusive outcome are relaxed compared to the case where consumers are long-run players. Because consumers do not behave strategically, we can rule out actions to break full collusion. Moreover, as consumers do not consider future rents, firms do not include the rents in the deviation price, which is now equal to $\tilde{p} = 1 - S$. In other words, the condition (IC_a) is harder to fulfil. As a result, full collusion is an equilibrium if and only if no firms wants to unilaterally break collusion.[18]

$$\frac{1}{1-\delta} \geq 2\left[(1-S)+\delta\frac{p(2)}{1-\delta}\right] \rightarrow \delta \geq \frac{-1+\sqrt{1-2S}+2S}{2S} = \underline{\delta}(S),$$

The case of short-sighted consumers is equivalent to the case with forward-looking consumers who do not put any weight to the future $\delta_{co}=0$. We denote by C_{SS} the set of parameters (δ, S) sustaining collusion. This region is represented in Figure 3.4.2 when $\delta_{co}=0$.

With short-sighted consumers, the feasibility of collusion does not depend on the number of consumers: if a firm attracts one consumer, it must attract all the others as well. Thus, our result easily extends to an arbitrary number of consumers (or a continuum). We obtain the following corollary.

Corollary 2. *In a duopsony, collusion is easier to sustain if consumers are short-sighted. The opposite holds for a continuum of consumers if these do not coordinate on the Pareto dominant outcome.*

3.6.2 Short-run Consumers

We proceed by studying how the feasibility of collusion change with the level of consumers' interactions with the firms. We examine the simplest case when consumers interact with firms only once. Here consumers die after purchase and at each period a new generation of consumers arrive and firms equally share the market. As in the case of short-sighted consumers, short-run consumers shop at the firm yielding the highest current utility, but their captivity does not depend on the past purchasing decisions, i.e. whatever happened previously, firms equally share the market at the beginning of each period.

Punishment Path

Because the captivity for short-run consumers is fixed, our game is an infinite repetition of the same stage game. As a result, the only possible state of the world at the beginning of each period is $\xi=1$. As consumers live for only one period, the punishment path coincides with the static Nash equilibrium $\mathbf{p}(\xi=1)=\mathbf{p}(N)$. Moreover, the discontinuity of the payoff function leads to the non-existence of a pure strategy equilibrium in prices whenever the level of switching costs is sufficiently low, i.e. whenever deviating from the collusive price is profitable. (See Dasgupta and Maskin [14] and Shilony [49].) This happens when the magnitude of switching costs is lower than one half the reservation price of consumers.

[18] In a companion note, we consider the case in which the firm does not attract the consumer of the rival in the moment of deviation. We show that the continuation game is one characterized with mixed strategies. We prove that such a strategy is never optimal. For low switching costs, any firm obtains strictly higher profits by attracting the consumer of the rival. For high switching costs the collusive outcome gives higher profits than any deviation.

Restricting attention to a symmetric equilibrium, the expected profit of a firm is:

$$\mathbb{E}\left[\pi(p)\right] = p \times \left[F(min(p+S,1)) - F(max(0,p-S))\right] + 2p \times \left[1 - F(min(p+S,1))\right],$$

where $F_i(p)$ is the cumulative distribution function of the competing firm. The first part of the expected profit represents the case where the difference in prices is below the level of switching costs and no consumer switches. The second part of the expected profit represents the case where the difference in prices is above the level of switching costs and consumers switch. Hence, in equilibrium the expected profit is:

$$\frac{V}{p} = 2 - \left[F(min(p+S,1)) + F(max(0,p-S))\right]. \tag{3.6.1}$$

Solving the previous expression in p and using the methodology proposed in Shilony [49], we introduce the following Lemma.

Lemma 3. *The Nash reversion equilibrium is characterized by:*

i) For $S \in \left(0, \frac{1}{2+\sqrt{2}}\right]$, the per period expected profit is $V = (1+\sqrt{2})\,S$ and consumers switch with positive probability.

ii) For $S \in \left(\frac{1}{2+\sqrt{2}}, \frac{1}{2}\right)$, two Nash equilibria (σ^I, σ^{II}) exist depending on whether indifferent consumers switch or not. In both equilibria, consumers switch with positive probability and the per period expected profits are:

(σ^I) *$V = 1 - S$ in an equilibrium when indifferent consumers switch.*

(σ^{II}) *$V = \frac{S+\sqrt{S(4+S)}}{2}$ in an equilibrium when indifferent consumers do not switch.*

iii) For $S \geq \frac{1}{2}$, the per period profit is $V = 1$ and no consumer switches.

Proof. This is a particular case of Shilony [49] when the number of firms is equal to two. The author shows that the support of prices lies in an interval not larger than twice the level of switching costs. However, with no specific assumption on the tie-breaking rule of consumers, we obtain two equilibria whenever $S \in \left(\frac{1}{2+\sqrt{2}}, \frac{1}{2}\right)$. We detail the formal proof in Appendix C.2. □

From Lemma 3, the expected profits of the firms are lower when indifferent consumers switch. When consumers are more prompt to switch, firms behave more aggressively which leads to lower expected profits.[19]

[19] We show in Appendix C.2 that in this case firms play with a positive probability the lower bound of the support. Conversely, when indifferent consumers do not switch, firms play the upper bound of the support with some positive probability, which coincides with the equilibrium obtained in Shilony [49].

Collusive Equilibrium

As with short-sighted consumers, short-run consumers do not make unilateral decisions to break the collusive outcome and we only need to verify that no firms wants to unilaterally deviate. The optimal deviation price is the same as when consumers are short-sighted, $\tilde{p} = 1 - S$, and the firms incentive constraint (IC_f) to sustain collusion is:

$$\frac{1}{1-\delta} \geq 2 \times (1-S) + \delta \frac{V}{1-\delta}. \qquad (IC_f)$$

Solving for the expected discounted value of the punishment path obtained in Lemma 3, we obtain the regions where collusion is possible.

Proposition 10.

1. *For $S = 0$, full collusion is an equilibrium for $\delta \geq \frac{1}{2}$.*
2. *For $S \in \left(0, \frac{1}{2+\sqrt{2}}\right]$, full collusion is an equilibrium for $\delta \geq \frac{1-2S}{2-(3+\sqrt{2})S} = \underline{\delta}(S)$.*
3. *For $S \in \left(\frac{1}{2+\sqrt{2}}, \frac{1}{2}\right)$, we have two possible equilibria:*
   - *If σ^I, full collusion is an equilibrium for $\delta \geq \frac{1-2S}{1-S} = \underline{\delta}(S)$ and partial collusion $\left(p^c = \frac{\delta - S(2-\delta)}{2\delta - 1}\right)$ for $S \in \left(\frac{1}{3}, \frac{1-\delta}{2-\delta}\right)$.*
   - *If σ^{II}, full collusion is an equilibrium for $\delta \geq \frac{2(1-2S)}{4(1-S)-(S+\sqrt{S(4+S)})} = \underline{\delta}(S)$.*
4. *For $S \geq \frac{1}{2}$, full collusion is always an equilibrium.*

Proof. We detail the proof in Appendix (C.2). □

Figures 3.6.1 and 3.6.2 recap the result of Proposition 10.

Full collusion is implemented whenever the value of switching costs is so high that attracting the consumers of the the rival is not profitable, i.e. $S \geq \frac{1}{2}$. Also, with relative small levels of switching costs, i.e. $S \in \left(0, \frac{1}{2+\sqrt{2}}\right]$, attracting the consumers of the rival is less costly and collusion is harder to implement. For the values of switching costs in $S \in \left(\frac{1}{2+\sqrt{2}}, \frac{1}{2}\right)$ the results depend on the behaviour of indifferent consumers. When indifferent consumers switch, the region sustaining full collusion is larger than when indifferent consumers remain loyal to their firm.[20] This result comes directly from the fact that

[20]For some values of parameters and switching indifferent consumers, only partial collusion is sustainable, i.e. the collusive price is strictly lower than 1. This range of parameters is depicted in red in Figure 3.6.1.

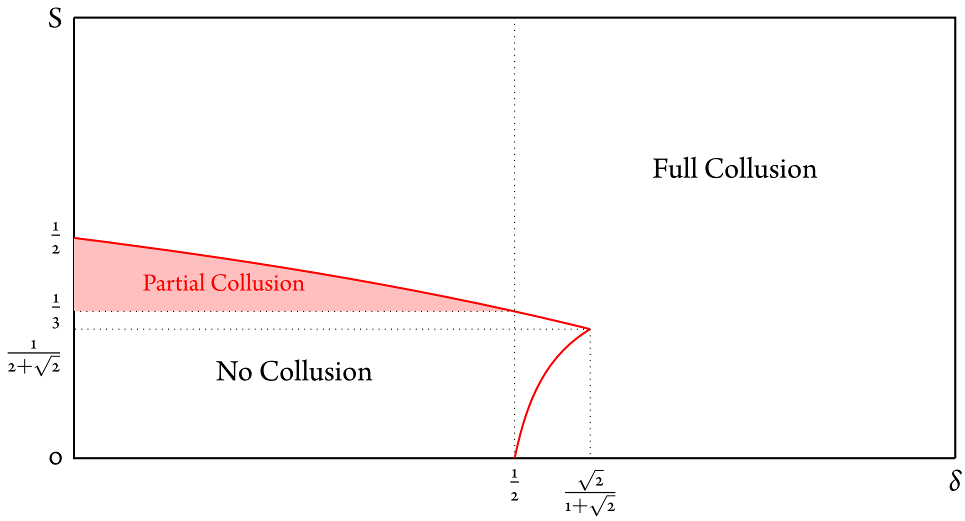

Figure 3.6.1: Collusion area with short-run consumers when consumers decide to switch when they are indifferent — The region in red represents a situation where collusion below full consumer's reservation price is sustained.

the punishment path after deviation is more severe in the case of switching indifferent consumers. These results extend to a continuum of consumers.

Remark 3. *With short-run consumers who switch when they are indifferent, collusion is harder to implement than with short-sighted ones.*[21] *As the immediate cost of deviation is the same, this comes from future gains from deviation being lower when consumers are short-sighted. Hence, although the deviating firm enjoys the whole market indefinitely when consumers are short-sighted, the incumbency price is so small in this case that a one time deviation in the short-run case is more profitable.*

[21] If σ^I is played, there is a small area on which collusion is possible for short-run consumers, but not short-sighted ones.

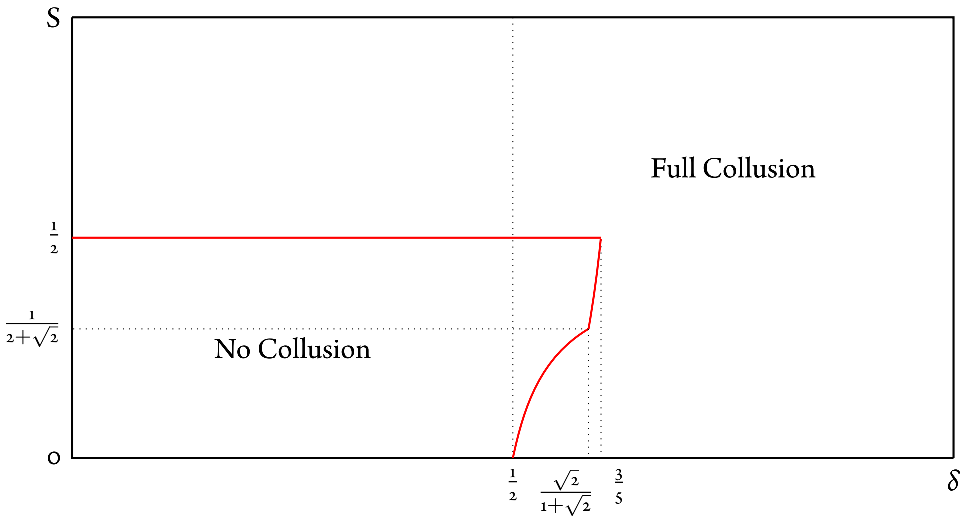

Figure 3.6.2: Collusion area with short-run consumers when consumers do not switch when they are indifferent.

3.7 Increasingly Loyal Consumers

Until now, we have considered that the consumers' switching costs are homogeneous. Moreover, switching costs are invariant over time and do not depend on the number of periods a consumer have purchased from a firm. Yet, homogeneity does not apply in many industries, and the level of switching costs can depend on the past purchasing history of a consumer. This occurs when there is learning associated to the usage of a product or in the presence of brand loyalty. In those markets, the costs associated with buying from another provider increase with the number of purchases of a single product. As a result, switching costs evolve endogenously and depend on the purchasing decisions of consumers. Here, we discuss the case in which switching costs increase with the number of periods a consumers remains captive. This captures situations of increasing synergies between sellers and buyers.

We denote the initial value of switching costs S_0 and it evolves over time according to a function $S_t = Z(t)$ that increases with the number of periods - t - a consumer remains captive to a firm. The possibility of heterogeneous switching costs arises as switchers have a lower switching costs than consumers that remain captive. Once a consumer has switched, his level of switching costs goes back to the initial level.[22] Because we aim at characterizing the pricing strategy of firms, we only consider the case

[22]Some recent studies, e.g. Biglaiser et al. [8], have considered the case where consumers differ in their switching costs. However, in those papers heterogeneity of switching costs is assumed, while heterogeneity here is considered endogenous.

of two long-run representative consumers that do not behave strategically.[23] To further simplify the analysis, we introduce the following assumption:

Assumption 6. *The function $Z(t)$ is bounded by the consumer's reservation price, i.e.* $\lim_{t\to\infty} Z(t) \leq 1$.

Therefore, the level of consumers switching costs always stays below the consumer's valuation for the product. The following Lemma introduces the equilibrium of our incumbency game.

Lemma 4. *The pair of prices that constitute an equilibrium of the incumbency game changes over time and is equal to:*

$$(p_t(\xi = 2), p_t(\xi = 0)) = (S_{t-1} - \delta(1-\delta)V(S,\delta), -\delta(1-\delta)V(S,\delta)) \quad (3.7.1)$$

where $V(S) = \sum_{t\geq 0} \delta^t \times Z(t)$ *is the present discounted value of switching costs. Moreover, no consumer switches in equilibrium.*

We detail the formal proof in Appendix (C.3). The result is fairly intuitive. As switching costs increases with captivity, so does the equilibrium prices in our incumbency game. At the beginning, and after a deviation from the collusive path have occurred, the price is low in order to retain consumers. This is similar to an "investing" phase. The objective of the firm is to increase the degree of captivity of his consumers.[24] Later, as switching costs increase, the firm increases her prices accordingly. This represents a "harvesting" phase where the firm obtain larger profits.

We proceed by studying to what extent increasing loyalty affect the sustainability of a collusive outcome. Because switching costs increase with captivity, firms want to deviate from collusion at the beginning by setting a deviation price $p^d = p(c) - S_0$. Therefore, the discounted profits from deviation are $\Pi^d = 2 \times (p(c) - S_0) + \delta\Pi\left(p(2)\right)$. The incumbency discounted profits are:

$$\begin{aligned}\delta\Pi\left(p(2)\right) &= 2\delta\left(S_0 - \delta(1-\delta)V(S,\delta)\right) + 2\delta^2\left(S_1 - \delta(1-\delta)V(S,\delta)\right) + \dots \\ &+ 2\delta^t\left(S_{t-1} - \delta(1-\delta)V(S,\delta)\right) + \dots \\ &= 2 \times \sum_{t\geq 0} \delta^{t+1}\left(S_t - \delta(1-\delta)V(S)\right) = 2\delta(1-\delta) \times V(S,\delta)\end{aligned}$$

Therefore, full collusion is an equilibrium of our game if the discounted profits from collusion are above

[23] We leave the case of forward-looking consumers and a continuum of them for future research.

[24] Observe that during this phase prices might be negative in order to keep the consumers captive. For some $t < \bar{t}$, we might have that $S_{t-1} < \delta(1-\delta)V(S)$.

the ones obtained from deviation.

$$\frac{1}{1-\delta} \geq 2 \times (1-S_0) + 2\delta(1-\delta) \times V(S,\delta) = \Pi^d$$
$$\iff 0 \geq 1 - 2\delta - 2(1-\delta) \times S_0 + 2\delta(1-\delta)^2 \times V(S,\delta) = Rhs(S_0,\delta)$$
$$\iff \delta \geq \delta^* \text{ where } \frac{2\delta^* - 1 + 2(1-\delta^*)S_0}{2\delta^*(1-\delta^*)} = (1-\delta^*)V(S,\delta^*)$$

and the value of the discount factor $\delta^* \in (0,1)$ that solves the equation is unique.[25] From this expression, collusion is easier to implement when the initial level of switching costs S_0 is large and harder when the discounted value of switching costs $V(S,\delta)$ is large. The former makes the deviation from collusion more costly and the latter leads to larger profits after deviation.

The following result establishes that for the same magnitude of switching costs, collusion is harder to sustain in the case of increasingly loyal consumers.

Proposition 11. *Assume $(1-\delta)V(S) = S$, then full collusion in the case of increasingly loyal consumers is an equilibrium if:*

$$\delta \geq \frac{-1 + S + \sqrt{1 + (S_0 - S)^2 - 2S_0} + S_0}{2S}$$

Compared to the case of homogeneous switching costs S, collusion is harder to sustain.

We only need to show that:

$$\frac{-1 + S + \sqrt{1 + (S_0 - S)^2 - 2S_0} + S_0}{2S} \geq \frac{-1 + \sqrt{1 - 2S} + 2S}{2S}$$
$$\iff \sqrt{1 - 2S_0 + (S_0 - S)^2} - \sqrt{1 - 2S} \geq S - S_0$$

which is always true as long as $S > S_0$ and $S > 0$.

The intuition of this result is simple. As the value of switching costs increases with captivity, and because consumers are short-sighted, attracting the consumers of the rivals is cheaper when consumers switching costs increase with captivity. Conversely, incumbency profits increases over time because consumers switching costs increase. Thus, whenever consumers have increasing switching costs, deviation from a collusive outcome is cheap and the punishment path after deviation is less severe. This two effects work in the same direction, making the implementability of a collusive outcome harder. We proceed the analysis by obtaining the regions of full collusion for a given function of switching costs.

[25] The uniqueness of δ^* comes from the right hand side decreasing with δ as $\lim_{\delta\to 0}(1-\delta)V(S,\delta) = S_0$ and $\lim_{\delta\to 1}(1-\delta)V(S,\delta) = 0$. And whenever $S_0 \leq \frac{1}{2}$, the left hand side is decreasing with δ as $\frac{\partial Lhs}{\partial\delta} = \frac{1}{2}\left(\frac{1}{(\delta-1)^2} + \frac{1-2S_0}{\delta^2}\right) > 0$ and the limits $\lim_{\delta\to 0} Lhs = -\infty$ and $\lim_{\delta\to 1} Lhs = +\infty$.

Example   Switching costs increase with the periods of captivity - t - according to the function:

$$Z(t) = \frac{b+1-a^{-t}}{b+1}; \qquad a \in (1,2) \text{ and } b \in \mathbb{R}_+$$

The parameter b affects the initial level of switching costs $S_0(b) = \frac{b}{b+1}$, and the parameter a stands for the intensity of increase of switching costs. The function is both increasing and concave with respect to the number of periods a consumer have stayed captive with the firm. By differentiating the function $Z(t)$ with respect to - t - we obtain:

$$\frac{\partial Z(t)}{\partial t} = \frac{a^{-t}\log(a)}{1+b} > 0; \qquad \frac{\partial^2 Z(t)}{\partial t} = -\frac{a^{-t}\log(a)^2}{1+b} < 0$$

To examine when a collusive outcome may be sustained in equilibrium, we only need to consider that no unilateral deviation by a firm takes place.[26] To facilitate the exposition, we introduce two measures, the discounted value of switching costs and the discounted profits from deviation:

$$V(S \mid b,a) = \sum_{t\geq 0} \delta^t \times Z(t) = S_0(b) + \sum_{t\geq 1} \delta^t \times \left(\frac{b+1-a^{-t}}{b+1}\right) = \frac{1}{B(1-\delta)}\left[b + \delta\left(1 + v^1(a,\delta)\right]\right)$$

$$\begin{aligned}\Pi^d(b,a) &= 2 \times [(1 - S_0 + \delta(1-\delta))V(S \mid b,a)] = 2 \times \left[\frac{1+\delta(b+\delta)}{b+1} + \frac{\delta^2}{b+1}v^1(a,\delta)\right] \\ &= \frac{2}{B}\left[1 + \delta^2\left(\frac{b+\delta}{\delta} + v^1(a,\delta)\right)\right]\end{aligned}$$

where $B = (b+1)$ and $v^1(a,\delta) = -\sum_{t\geq 1}\delta^t \times a^{-t}$ is the discounted value of the intensity of growth of switching costs.

We see that the discounted deviation profits increase and decrease with both the intensity of growth of the captivity of consumers and the initial level of switching costs:

$$\frac{\partial \Pi^d(b,a)}{\partial b} = -\frac{2}{B}\left[1 + \delta^2\left(\frac{b+\delta}{\delta} + v^1(a)\right)\right] + \frac{2\delta}{B} = -\frac{2}{B}\left[1 - \delta + \delta^2\left(1 + v^1(a,\delta)\right)\right] < 0$$

$$\frac{\partial \Pi^d(b,a)}{\partial a} = \frac{\delta^2}{B}\frac{\partial v^1(a,\delta)}{\partial a} > 0$$

This implies that the likelihood of sustaining a full collusive outcome is affected by both parameters. The results are represented in Figure 3.7.1. We observe that the necessary level of firms' discount factor to

[26] In this example, deviations by consumers are not considered since they do not behave strategically.

sustain collusion increases with the intensity of growth of switching costs. On the other hand, the initial level of switching costs and the cost of deviation also increase with the parameter b, making full collusion easier to sustain.

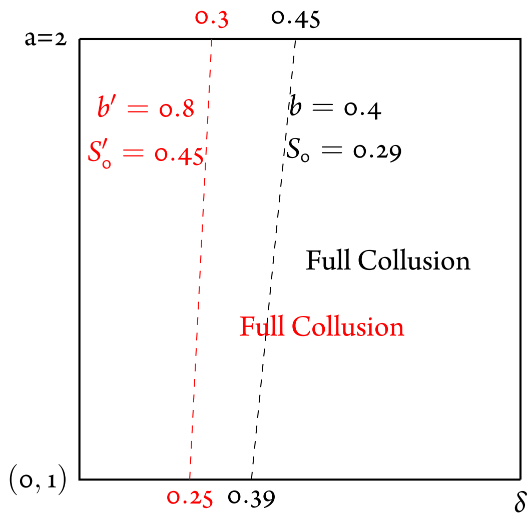

Figure 3.7.1: Area of full collusion with increasing switching costs as a function of the initial switching cost and the intensity of increase.

Therefore, the fact that switching costs might be evolving with respect to the purchasing history of consumers also has an important effect on the feasibility of collusion. Here, we have established that collusion is harder to sustain whenever switching costs increase with the number of periods a customer have remained captive to a firm. Hence, a competition authority should pay more attention in industries where the level of switching costs remains constant. However, we have not tackled the issues of increasingly loyal and forward-looking consumers, which is left for future research.

3.8 Conclusion

We have seen that the sustainability of collusion depends on the fundamentals of the economy as well as on the type of consumers on the market. Their level of sophistication affects the incentives to deviate from the collusive path while the frequency with which they interact with the firms affects the severity of punishments. Hence, in the presence of two representative consumers, maintaining a fully collusive outcome is more difficult when consumers are forward-looking. But this result does not extend to a continuum of forward-looking consumers. When consumers build loyalty over time, the incumbency game is characterized by well defined *investing* and *harvesting* phases. In this case, for a comparable level of switching costs, collusion is harder to sustain. Overall, if they could, firms would have incentives to set large switching costs at the beginning to facilitate collusion.

A competition authority should therefore take into consideration the level of consumers' sophistication as well as their level of interaction with firms if switching costs are present. Furthermore, the number of consumers and their beliefs about the behaviour of their peers seem to have a great impact on the sustainability of a collusive outcome. Educating consumers on the effects of their current purchasing decisions on future equilibrium prices and on the benefits of coordination seems a sound strategy to reduce the possibility of tacit collusion.

4

A Welfare Assessment of Revenue Management Systems

co-authored with Marc Ivaldi and Jerome Pouyet

Abstract

We study the welfare impact of revenue management, i.e. intertemporal price discrimination when the product availability is limited both in time and quantity, and consumers' arrival is random. This practice is particularly relevant, and widely spread, in the transport industry, but little is known about its implications on profits and consumer surplus. We develop a theoretical model of revenue management allowing for heterogeneity in product characteristics, capacity constraints, consumer preferences, and probabilities of arrival. We also introduce dynamic competition between revenue managers. We solve this model computationally and recover the optimal pricing strategies. We find that revenue management is welfare enhancing. Revenue managers face two types of constraints: a limited booking period and fixed capacities. Previous sales affect the relative slackness of these two constraints, explaining price variations. Profits increase as the practice offers more leeway to the seller compared to posting a fixed price throughout the booking period. Total consumer surplus also increases for a wide range of specifications, as revenue management raises the number of sales. In the presence of heterogeneous consumers, consumers with low price sensitivity subsidize ones with high price sensitivity when demand is low but both types benefit from the practice when demand is high. This sheds some light on the impact of revenue management on the surplus of business and leisure passengers.

4.1 Introduction

We study the welfare impact of revenue management and the dynamics of fares resulting from this pricing strategy. We find that revenue management is generally welfare-enhancing. Revenue management is widely used in transports, especially by airline and railway companies, but also to manage hotel bookings and ticket sales. This form of intertemporal price discrimination allows the seller to manipulate the price of a product whose availability is constrained both in time and in quantity. In a framework in which consumer arrival is uncertain, the seller can react to the demand realization: she increases or decreases her prices given the quantity already sold and the remaining time in which to sell the residual capacity.

Whether revenue management is welfare-enhancing or not is a difficult question which remains, to the extent of our knowledge, unanswered. The practice might seem unfair, and on the ground that it increases firms' profits, raises suspicions that it harms consumers. However, forcing the firm to post a fixed price does not ensure that it will post an affordable one. When restricted to a fixed price, the revenue manager can no longer reward early bookers through rebates or propose last minute deals at the end of the booking period.

We develop a model of revenue management allowing the seller to simultaneously propose different products, heterogeneous in their characteristics but also in their availability, to different types of consumers. For instance, a revenue manager in a railway company could propose several trains on a given day for the same origin-destination leg. These trains could differ in their departure time, number of stops along the way, or number of seats. Each is characterized by a capacity constraint, i.e. its number of remaining seats. As a revenue manager is also likely to react to the prices posted by her competitors, we introduce dynamic competition between revenue managers.

We computationally solve this dynamic program by simulating stochastic arrivals of consumers and their purchasing decisions. We then compute the subsequent optimal reaction of the revenue manager. Our results are robust to several specifications for the arrival of consumers as we allow for constant and increasing arrival rates.

Our approach allows us to answer the difficult questions of finding the equilibrium pricing strategies and assessing the welfare impact of revenue management despite the absence of a closed-form solution to our model. Another method to understand the strategies played at equilibrium and to carry out this welfare analysis would have been to estimate a structural model of revenue management. It is however difficult to procure sufficiently detailed data on this topic as they are a very sensitive and strategic piece of information.

We first consider a simple model of revenue management with homogeneous products and only one type of consumers. The revenue manager only faces one capacity constraint and one time constraint. This allows us to illustrate the forces at play in revenue management: the intensity of the constraints faced by the revenue manager can be represented by the ratio of remaining units over the remaining time to sell them. For a given number of units, if the remaining time to sell them decreases, the ratio increases, and the revenue manager is pressured into lowering her prices. If, on the other hand, for a given deadline, the revenue manager has less units to sell, the unit-time constraint faced by the revenue manager is relaxed and she will try to sell the last units at a higher price.

We are able to recover a standard feature of the industry pricing: if the arrival rate of consumers is sufficiently high compared to the total capacity, prices start low and increase as time goes by. At the very end of the booking period, prices can dramatically fall as the manager does not want to keep unsold units. When the intensity of demand is low compared to the total capacity, the probability of the revenue manager being constrained by the number of available units is almost zero and intertemporal price discrimination does not occur in practice. We also show that compared to an optimal fixed price, revenue management strongly increases consumer surplus through a higher number of sales without hurting the profits of the company.

We also consider the more complex and realistic case in which heterogeneous consumers have to choose between heterogeneous products, for instance an off-peak and a rush-hour train. The revenue manager now faces one time constraint but several capacity constraints. Solving for the optimal pricing strategy of the revenue manager and for the optimal fixed price, we compare the producer and consumer surpluses between these two pricing strategies. For low intensities of demand compared to the total capacity, revenue management allows price-sensitive consumers to be subsidized by the ones with a lower price elasticity. However both categories of consumers benefit from the practice as the intensity of demand increases for a given total capacity. Applied to the transport industry, this provides some insights about the impact of revenue management on leisure and business passengers.

In the third part of the paper, we introduce both indirect and direct competition in the model, and study their impact in terms of profit and welfare. We find that compared to the monopoly case and holding demand constant, direct competition substantially increases load rates through lower prices. We also study the impact of ex-ante strategic decisions on social welfare, such as allowing revenue managers to choose the total capacity they propose before they compete against each other. In this context, we find that competition between two revenue managers reduces producer surplus only slightly compared to the monopoly case but allows consumer surplus to be twice as large.

We review the literature in section 4.2 and present the model and our algorithm in sections 4.3 and 4.4. Results for the simplest one-product homogeneous-consumers case are presented in section 4.5. Heterogeneous consumers are introduced in section 4.6. We analyse dynamic competition between revenue managers and ex-ante competition in capacities in sections 4.7 and 4.8.

4.2 Literature Review

The literature on revenue management is wide both in operations research and economic theory. Our paper fills the gap between the two fields, answering a truly economic problematic in the realistic framework of operations research.

In operations research, Talluri and Van Ryzin [51] and Talluri and Van Ryzin [52] introduce a choice-based model of revenue management and show the two forces driving the manager's actions: the will to boost demand by selling at a low price, and the incentive to post high prices after a sufficient number of sales.[1] The first force is prevalent at the beginning and at the end of the booking period. They also show that the choice-based approach gives better results than other methods of revenue management. We show that some of their results do not hold when we allow the capacity constraint of the revenue manager to be multidimensional, e.g. when the revenue manager maximizes profits over two different types of products instead of one. In that case, the revenue manager may be willing to post prices which are not efficient according to their definition. McAfee and Velde [35] study a similar question, restricting to demand functions satisfying constant price elasticity, which allows them to derive a closed-form solution. They also use a reduced-form approach by modelling arrival of consumers and their purchasing decision as a simultaneous process. To the extent of our knowledge, Vulcano et al. [55] is the first attempt to estimate consumers preferences and to carry out counterfactuals in a revenue management context. Gallego and Van Ryzin [22] introduce the notion of multiple products and associated limited resources although they adopt a reduced-form approach for the consumers' behaviour. Modarres M. [37] study the pricing strategy when cancellation is possible and revenue managers do not update their prices every period. McGill and Van Ryzin [36] provide a comprehensive overview of the issues raised by revenue management. This operations research literature does not provide any welfare analysis of revenue management, but rather focuses on the best way to implement it from the perspective of the seller. Despite some similarities in the way we model the maximization program of the revenue manager, our aim is radically different. We also extend these models by introducing strategic choices and strategic interactions from and among revenue managers.

The economic literature, starting with Stokey [50], has dealt with intertemporal price discrimination but not really with revenue management from an operational perspective. For instance, Dana Jr [13] studies how a firm can use revenue management to smooth demand peaks and reduce capacity costs. He models revenue management as price discrimination between two flights differentiated by their departure time, but does not consider fluctuating prices for a given flight. Related to our topic, although more theoretical, Gershkov and Moldovanu [24] use a mechanism design approach to solve the revenue

[1] A choice-based model is one which explicitly models consumers' choices.

management problem. This method has the advantage of providing a model with a closed-form solution. However it stands on a theoretical ground which is far from the practitioners' world. Hörner and Samuelson [27] study the allocation of a good when consumers are forward-looking. This framework certainly improves the theory of revenue management but is, in our opinion, less useful to study more practical issues such as its welfare impact. Indeed, forward-looking consumers are supposed to be able to observe and react to any price change from the revenue manager from the moment they realize they need a ticket to the end of the booking period. We feel that search costs or aversion towards the risk of not finding a seat might also play a role when booking a ticket. In that case, even forward-looking consumers could behave like impatient ones. In practice, in the transport industry, we feel that consumers looking for a ticket make their purchasing decisions comparing prices of travelling options with close departure dates rather than trying to anticipate future prices. Because of this, we choose to simplify the analysis and model our customers as impatient, which is a standard feature in the operations research literature. This allows us to model a more realistic framework including several substitutable products to choose from. We also address the possibility of partially forward-looking consumers in one specification of our model by allowing the willingness to purchase the products to increase over time. This represents a situation in which consumers are aware that there might not be any products left if they wait but cannot anticipate the pricing strategy of the revenue manager. Lazarev [32] empirically studies the welfare impact of intertemporal price discrimination. His approach however does not explicitly take into account revenue management but assumes an exogenous dynamics of fares. Williams [57] estimates a dynamic model of revenue management and price discrimination using a framework similar to ours. He finds that revenue management benefits consumers on average, which confirms our first result. However, he avoids the problem of substitutability between products by focusing on single flights. We also conduct an analysis of revenue management when the revenue manager can propose several types of products.

Finally, Goodwin [25] and Wardman and Shires [56] both conduct reviews of price elasticities' estimates in transports, respectively with a focus on short-term vs. long-term and on the types of consumers. To carry out our analysis, we choose values of parameters giving us price elasticities consistent with the average estimates in this literature.

4.3 The Model

4.3.1 Notations

Our revenue management problem features a seller, the revenue manager, proposing several products, constrained in quantity, to a set of potential consumers. The revenue manager can only sell her products during a finite period of time, which corresponds to the booking period in the transport industry. In our model, this period starts at date T and ends at date 1. Consumers who purchased the products consume them at date 0.

We allow products to be heterogeneous in their characteristics and different types of products can be simultaneously proposed to consumers by the revenue manager. For instance, a revenue manager can simultaneously propose a first-class ticket and a second-class ticket. All products are constrained in quantity and each type of products has its own constraint. In the previous example, the number of first class tickets that the revenue manager can sell is constrained by the number of first-class seats. Products characteristics are denoted by $i \in I$.[2] X_i denotes the remaining capacity for attribute i. The vector of remaining capacities for all constrained characteristics in I is denoted X.

The range of possible prices for a product with attributes i is denoted P_i and contains a price equal to $+\infty$ for all i. At each period of the booking process, the revenue manager chooses the menu of tariffs she proposes to the consumers. $\tilde{p} \in \prod_i P_i$ is the vector of selected tariff options for all products. We assume that when $\tilde{p}_i = +\infty$ is chosen, the product is sold with probability 0. This tariff option is therefore equivalent to closing the sales of the product. When bought, a product with attributes i is exchanged at price $p_i = \tilde{p}_i$.

We assume that no more than one consumer can buy a product within a period t of the booking process.[3] At each period $t \in \{T, \ldots, 1\}$, a consumer arrives with probability λ_t, observes the available tariffs, and chooses whether or not to purchase. The dummy d_i takes value 1 when a product with attribute i is purchased. When the consumer selects the outside option, $d_0 = 1 - \sum_i d_i = 1$.

In the following, we sometimes use the expressions *market size* or *intensity of demand* when referring to the average rate of arrival $\mathbb{E}(\lambda_t)$.

4.3.2 The Bellman Equation

At each period t of the booking process, the revenue manager observes the remaining capacity X and chooses a menu of tariffs so as to maximize her overall profit. If she sells a product during the period, she gets the instantaneous profit given by the price of the product plus her continuation value when there is one less unit of product to sell. If she does not sell, she just gets her continuation value for a remaining capacity X. Selling occurs if a consumer arrives and decides to buy. Not selling occurs either because no consumer arrived or the consumer chose not to buy. We denote $Pr(d_i = 1|\tilde{p})$ the probability that a consumer buys a product i conditional on a proposed menu of tariffs $\tilde{p}$. We shall define this probability later when we explicitly model the consumers' behaviour.

[2]For the sake of simplicity, we chose to model only the products characteristics which are physically constrained. We could easily add another layer of differentiation, such as the flexibility of a ticket. But this would not change the flavour of our results.

[3]Although this might seem a bit restrictive, we can choose T as large as needed to ensure that this assumption is empirically justified.

Hence, the choices of the revenue manager must satisfy the following Bellman equation:

$$V_t(X) = \max_{\tilde{p}\in\Pi_i P_i} \sum_i \lambda_t Pr(d_i = 1|\tilde{p}) \left[\tilde{p}_i + V_{t-1}(X_i - 1, X_{-i})\right] + \left[1 - \lambda_t + \lambda_t Pr\left(d_0 = 1|\tilde{p}\right)\right] V_{t-1}(X) \tag{4.3.1}$$

subject to the final constraints:

$$\begin{cases} V_0(X_i) = 0 & , \forall X_i \\ V_t(X) = 0 & , \text{ whenever } \sum_i X_i = 0 \\ \tilde{p}_i = +\infty & , \text{ if } X_i = 0 \end{cases} \tag{4.3.2}$$

Denoting $\mathbb{E}(\pi|\tilde{p})$ the instantaneous expected revenue given $\tilde{p}$, we can rewrite Equation 4.3.1 as:

$$V_t(X) = \max_{\tilde{p}\in\Pi_i P_i} \lambda_t \left[\mathbb{E}(\pi|\tilde{p}) - Pr(d_0 = 0|\tilde{p})V_{t-1}(X) + \sum_i Pr(d_i = 1|\tilde{p})V_{t-1}(X_i - 1, X_{-i})\right] + V_{t-1}(X) \tag{4.3.3}$$

subject to the final constraints 4.3.2.

In the last equation, $Pr(d_0 = 0|\tilde{p})$ represents the total probability of purchase if the menu $\tilde{p}$ of tariffs is chosen.

The three final constraints are respectively given by the fact that sales are closed at the end of the booking process, the limited capacities, and the fact that the revenue manager has to close the sales of a type of products when it is no longer available.

4.3.3 Demand Modelling

The probability with which a consumer buys a product with characteristics i given $\tilde{p}$ is determined by the way we model demand. We adopt here a multinomial logit approach, in which the utility of a consumer arriving in period t is given by:

$$\begin{cases} u^t(p_i) = v + a u_i - \gamma p_i + \varepsilon_{it} & , \text{ when } d_i = 1 \\ u^t_0 = \varepsilon_{0t} & , \text{ when the outside option is chosen.} \end{cases} \tag{4.3.4}$$

In the above equation:

- ε is the random part of the utility and is drawn from a type-I extreme value distribution, with location parameter μ equal to the Euler-Mascheroni constant (≈ 0.5772) and scale parameter $\sigma = 1$. These parameters ensure that ε is a zero-mean random variable with a variance equal to $\frac{\pi^2}{6}$. Those noises are i.i.d across options and consumers.

- $\bar{u}_0 = 0$ in case of no purchase, and $\bar{u}(p_i) = v + au_i - \gamma p_i$ in case of a purchase of a product with attributes i at price p_i, are the deterministic parts of the utility. u_i represents the observable characteristics of product i. The parameters a and γ respectively represent the consumers' sensitivity to these characteristics (such as the comfort class) and to the price. v is a scale parameter. Note that we can allow these parameters do depend on consumers' observable characteristics.

Following Ben-Akiva and Lerman [6], the choice probabilities $Pr(d_i = 1|\tilde{p})$ are given by:

$$\begin{cases} Pr(d_i = 1|\tilde{p}) = \frac{e^{\bar{u}(p_i)}}{\sum_{j\in I} e^{\bar{u}(\tilde{p}_j)} + 1} \\ Pr(d_0 = 1|\tilde{p}) = \frac{1}{\sum_{j\in I} e^{\bar{u}(\tilde{p}_j)} + 1} \end{cases} \tag{4.3.5}$$

Notice that the final constraints of the Bellman equation imply that the revenue manager posts an infinite price p_i when a product i is sold out. Looking at Equation 4.3.5 above, it means that the probability of purchasing such a good is zero. If all products are sold out, all prices are set to infinity and the probability of choosing the outside option is 1.

Remark 4. *The relationship between γ and the price elasticity of demand $\eta(\tilde{p})$ is easily derived. Denoting $S_i(\tilde{p})$ the market share of a product with attributes i when consumers face a price vector $\tilde{p}$, we get:*

$$\eta(\tilde{p}) = \frac{\partial S_i}{\partial p_i}\frac{p_i}{S_i} = -\gamma p_i(1 - S_i) \tag{4.3.6}$$

4.3.4 Comments on the Theoretical Model

Compared to other models of revenue management in the operations research literature, we differentiate ourselves in the following ways:

- We allow for a general product set, described by several capacity constraints.
- We choose a flexible structural approach to model purchasing decisions.
- The revenue manager can choose the price of each available product but also which products to propose using the option of an infinite price. By comparison, some revenue management papers assume that if two products are identical in all their characteristics but their prices, they constitute two different products. This leads to unrealistic consumer behaviours: when proposed

simultaneously to the low price product, the high price product can be chosen by consumers with some positive probability.

Our model is therefore flexible enough to accommodate realistic pricing and consumption decisions. This flexibility has a cost: our model does not have any closed-form solution and some of the theoretical results that can be found in the literature do not extend here. More details about these issues can be found in Appendix D.1.

The best approach to solve our model is to use computational techniques.

4.4 Algorithm

To recover the optimal pricing strategy of the revenue manager and assess the welfare impact of revenue management, we construct an algorithm that solves the revenue manager optimization problem and finds her optimal action at any given period during the booking process and for any value of the state variable X.

First, we set the values of the different parameters driving consumers' choices as well as the choice sets of the revenue manager, $P_i \, \forall i$. All these values are chosen so as to be as realistic as possible. Values of parameters for the preferences of the consumers were chosen to match price elasticities estimated in the economic literature. (See Goodwin [25] and Wardman and Shires [56].) The number of tickets that the revenue manager has to sell matches the average capacity of a plane or a train. In our computations we choose capacities going from 200 seats up to 400 seats. The price grid was constructed based on observations about prices of train tickets in France and intra-European flights.

Given parameters for the preferences of the consumers and product characteristics, we are able to construct the purchase probabilities for each product as well as the expected instantaneous revenue using the demand model described above.

We therefore have all the elements necessary to solve the Bellman equation faced by the revenue manager. We start from the final constraints where the continuation value of the Bellman equation is zero and we compute the value functions at previous dates and for all possible states by iteration. More precisely, we fill a matrix with $T+1$ rows and $\sum_i X_i$ columns. The first column corresponds to the case in which all products are sold and is therefore filled with zeros. The first row corresponds to the end of the booking period when the revenue manager can no longer sell the products. It is also filled with zeros. We can find the value functions in the second row because we know all values in the preceding row are zero. Using the Bellman equation, the value functions in the second row are the maximum expected revenue conditional on the available products. A maximization program running on the finite set of prices available to the revenue manager gives us this value function as well as the optimal menu of prices. The continuation values in period 1 (second row) are plugged into the Bellman equation, and we can solve it for period 2. We continue this iterative process until period T. This gives us the optimal actions of the revenue manager for each t and X.

Finally, we simulate arrivals of potential consumers during the booking period by random draws from a binomial distribution with parameter λ. We also generate the random shocks ε_{it} in the utility of the consumers and recover their purchasing decisions and surplus. Profits and total consumer surplus are computed as the sum of individual purchases and surplus. This being a stochastic process, we repeat the whole operation hundreds of times so as to recover the expectation and variance of the profit and consumer surplus.[4]

4.5 The Distributive Properties of Revenue Management

Although discrimination often creates an inefficient allocation of goods among consumers, it can sometimes increase the consumer surplus through a higher volume of sales. In this section we show that this is indeed the case with revenue management, using the model developed in section 4.3, and realistic values of parameters.

We develop here the simplest case possible with homogeneous products and only one type of consumers. Products are only differentiated through their price. Average values of proposed prices, profit and consumer surplus are found using the algorithm described above. The values of parameters we chose are summarized in Table D.2.1 of Appendix D.2, column (1). We simulate this setup for constant arrival rates:

$$\lambda_t = \lambda, \ \forall t \text{ with } \lambda \in \{0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9\}$$

4.5.1 Price Variations in Revenue Management

In a first setup, we illustrate the revenue management practice using a range of possible prices, P, equal to $\{40, 90\}$.[5] The selected values of parameters and Equation 4.3.6 allow us to compute the price elasticity of demand in this case: $\eta \approx -0.51$ when $p = 40$ and $\eta \approx -2.1$ when $p = 90$. These are in line with price elasticities' estimates in the rail industry: in Goodwin [25], the average price elasticity for the demand for train tickets is -0.79 over 92 quoted values. In a more recent review, Wardman and Shires [56] find an average price elasticity of -0.9 over 456 inter-urban rail demand elasticities in Great Britain, with highest elasticities around -3.2.

[4] In the remaining of the paper, we present average profit and consumer surplus across all the simulations. Prices at each period of the booking process are also averaged across all simulations. The number of simulations varies between 500 and 1000 depending on the specification. (See Table D.2.1 of Appendix D.2.)

[5] These bounds have been chosen after observing the change in ticket prices on the second class Paris-Lyon over the whole booking period.

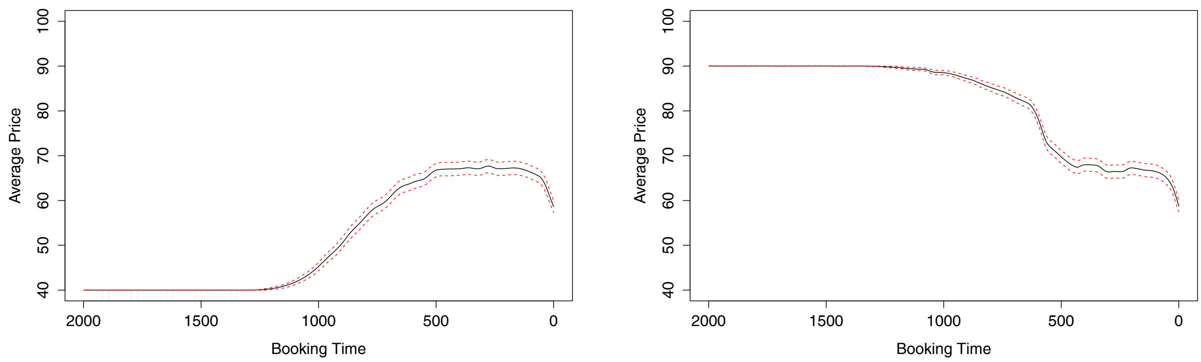

(a) (b)

Figure 4.5.1: Change in posted average prices as a function of the booking period. Figure (a) deals with a demand of medium intensity, $\lambda = 0.4$; Figure (b) deals with a demand of high intensity, $\lambda = 0.7$. Note: The averaging is performed over all the simulations at a given period of the booking process. To facilitate the reading, we use a non-parametric fit of average prices over the booking period. The confidence interval in red is the smoothed confidence interval of the mean.

Figure 4.5.1 represents the change in posted average prices in such a setup for $\lambda = 0.4$ and 0.7. The averaging is performed over all the simulations at a given period of the booking process. We notice that for a medium market size, the revenue manager tends to start with a low price: the time constraint dominates the capacity constraint. In that case, the revenue manager has the incentive to sell fast as many units as possible. When enough sales are made, prices start to increase: the capacity constraint now dominates and the incentive to sell the remaining units at high prices counterbalances the incentive to sell all units. At the end of the booking period, the opportunity cost of having unsold units dominates again and prices go down.

When the market size is larger compared to the total capacity, the frequent arrival of consumers makes it very likely that all units will be sold by the end of the booking period. Therefore, the capacity constraint already dominates at the beginning of the booking period and prices start high. This explains the shape of the price curve when $\lambda = 0.7$.[6]

Note that this general shape is not a feature of the particularly simple price structure we adopted in this example. We performed the same exercise with five different prices between 20 and 100 and reported the

[6]In practice, prices vary through time according to shape (a) rather than shape (b). This would mean that the proposed capacity is quite high compared to the market size. A possible explanation is that given the risk stemming from the random arrival of consumers, revenue managers prefer to underestimate the market size when maximizing their profit. However, such risk aversion is not modelled here. Another explanation would be that firms choose not to reduce the proposed capacity so as to avoid a shortage of units to sell and possible consumer discontent.

price variations for $\lambda = 0.5$ in Figure D.8.1 of Appendix D.8.1. Although variations are less pronounced, the general shape remains the same.

This shape of the price curve is standard in the revenue management literature: it corresponds to the price curve found in McAfee and Velde [35] and the incentives faced by revenue managers are mentioned in Talluri and Van Ryzin [51]. We can also retrieve this shape looking at transport prices, for instance airline prices, which usually increase as time goes by but can drop at the last minute before the departure of the flight.

4.5.2 Revenue Management vs. Optimal Fixed Price

Revenue management offers more leeway to the seller than an optimal fixed price strategy. It is therefore expected that a revenue management strategy including the optimal fixed price in the choice set of the seller must (at least weakly) increase her profit compared to the optimal fixed price alone. In this section, we test this prediction and measure to which extent revenue management including the optimal fixed price raises profits and affects consumer surplus.

To do so, we compute the optimal fixed price p^o and the average profit and consumer surplus associated with this price. We then construct the choice set P of the revenue manager to include p^o and some small variations around this price:

$$P = \{p^o - \xi, p^o, p^o + \xi\} \;\; \xi = 1, 2, \ldots$$

Although the revenue manager's choice set is small, Figures D.8.2 and D.8.3 of Appendix D.8.2 show that revenue management weakly increases both profits and consumer surplus. The increase is significant only for market sizes which are high compared to the capacity constraint. When $\lambda \geq 0.7$, revenue management increases profits by 1% compared to the fixed price strategy, and for the same values, the increase in consumer surplus lies between 2% and 3%. These values are significant and indicate that simply adding some leeway in the choice set of the seller can benefit both the company and the consumers.

We find no significant profit increase for $\lambda \leq 0.4$. In these cases, the price posted by the revenue manager is always the optimal fixed price. For these values of the arrival rate, the observation of a high volume of sales during the booking process is very unlikely. The revenue manager does not have any incentive to modify her posted price because she anticipates that the capacity constraint will not be binding. Further discussion of these results can be found in Appendix D.3.

4.5.3 Revenue Management and Increasingly Impatient Consumers

Consumers are obliged to choose the outside option if all products are sold out. Although we do not model the consumers' behaviour in a dynamic way, one way to introduce concerns for the future would be to consider consumers with an increasing willingness-to-pay.

Indeed, let us assume a semi-sophisticated consumer who is well aware of the possibility that she might not get a ticket if she waits but cannot fully anticipate the pricing strategy of the revenue manager. If this consumer arrives at the beginning of the booking period, she might think that she has enough time to find another option if she chooses not to buy immediately. Her value for the product is pretty low relative to her value for the outside option. But if she arrives at the end of the booking period, there is a high probability that all products will be sold out soon and she is running out of time to find a decent outside option. Her willingness to buy the product becomes high compared to the outside option.

Taking the example of someone looking for plane tickets, she might feel that she has plenty of other opportunities at the beginning of the booking period and might refuse a high-price ticket. At the end of the booking period however, her other options may be less enticing or riskier — for instance other means of transportation might not be available anymore. In this case, she might be willing to pay a higher price for her ticket.

We model this situation by allowing the scale parameter in the utility function of the consumer to be a decreasing function of $t \in [T, 1]$. Hence, when purchasing product i, the consumer gets the following utility:

$$u^t(p_i) = \nu + 1 - \frac{t-1}{T} + au_i - \gamma p_i + \varepsilon_{it} \qquad (4.5.1)$$

Revenue management might allow the seller to extract even more profit compared to committing to a fixed price strategy. Indeed, the seller can be flexible and adapt her prices to the arrival of consumers who are more willing to pay high prices.

As we want to compare this situation to the case in which consumers have always the same willingness to pay, we choose ν in Equation 4.5.1 to be equal to 1, thus ensuring that the expected utility for purchasing a product at price p_i averaged over the booking period is equal to their expected utility when they have a constant willingness to purchase.

Figures D.8.10 and D.8.11 of Appendix D.8.2 show our results when consumers have an increasing willingness to purchase the product. For $\lambda \leq 0.4$, revenue management has no effect on profits. When $\lambda \geq 0.5$, the relative impact of revenue management compared to a fixed price on profits is slightly stronger if consumers have an increasing willingness to purchase. Indeed, the increase in profits due to revenue management lies around 1%, which is similar to the case in which consumers have a constant willingness to purchase throughout the booking process. This impact is even stronger in some instances with a peak at 1.5% when $\lambda = 0.5$.

However, the effect of revenue management on consumer surplus in this situation becomes ambiguous as it increases or decreases consumer surplus depending on the value of the arrival rate. More surprisingly, revenue management has some positive effect on the consumer surplus even for low values of λ: as consumers now have an increasing valuation for the product, the revenue manager has an incentive to deviate from the optimal fixed price when $\lambda \leq 0.4$. When $\lambda \geq 0.6$, the relative impact of revenue management is similar in both cases of constant and increasing willingness to purchase, around $2 - 3\%$. But it appears that for intermediate values of the arrival rate, the impact of revenue management on consumer surplus becomes negative. This demonstrates that in this situation, the fact that revenue management creates surplus through a higher number of sales can be offset by a greater ability to capture consumer surplus. This is confirmed by the observation that the value for the arrival rate yielding the highest profits is also the one yielding the lowest consumer surplus.

4.5.4 Revenue Management with Noisy Arrival Rates

Another practical issue is the fact that the revenue manager can face some noise in her prediction of the arrival rate of consumers. Even though the prediction is good on average, i.e. over several similar booking periods, for a given booking period unobserved shocks might shift this arrival rate. Formally:

$$\tilde{\lambda} = \lambda + \nu$$

where $\tilde{\lambda}$ is the realized arrival rate and ν is a Gaussian white noise.

Revenue management has the extra advantage over a fixed price to be flexible enough to limit the loss of profits due to these unobserved shifts. Suppose $\tilde{\lambda} > \lambda$. The seller underestimates the size of the market and might post a price which is too low. However, because she sells her products more quickly, she will be able to raise her prices on her remaining inventory over a longer period of time.

To model this situation, we use the same framework as in section 4.5.2 except that we add a normally distributed noise to λ when we simulate the arrival of consumers.[7] We choose $\nu \sim N(0, 0.02)$. Hence, the optimal pricing strategy of the revenue manager is computed using λ but the actual arrival rate of consumers is simulated using $\tilde{\lambda}$. Other values of parameters remain unchanged.

Figures D.8.12 and D.8.13 of Appendix D.8.2 show our results when the arrival rate of consumers is noisy. For $\lambda \leq 0.4$, the realized arrival rate of consumers is too low for revenue management to make any difference with an optimal fixed price. In the following, we focus on $\lambda \geq 0.5$. Clearly, the relative impact of revenue management compared to a fixed price on profits is stronger when the arrival rate is noisy.

[7] The added noise is constant for the whole booking process.

Indeed, revenue management increases profits from 1 to 2% whereas the increase was at most 1% in the absence of noise. The fact that the seller cannot perfectly infer the arrival rate of consumers gives therefore an additional advantage to revenue management compared to a fixed price.

However, the impact on consumer surplus is ambiguous. In some cases, e.g. $\lambda = 0.5$, revenue management lowers consumer surplus by 1 to 2% whereas there is no significant impact in the noiseless case. For higher values of λ, as revenue management leads to more purchases, consumer surplus still increases compared to a fixed price, but to a lesser extent than in the noiseless case. Our interpretation is that when arrival rates are noisy, the sub-optimality of the revenue manager's pricing strategy given $\tilde{\lambda}$ creates inefficiencies. For instance, if the arrival rate is overestimated, prices will be too high for too long at the beginning, and this will preclude consumers to purchase. This phenomenon is not offset by cases in which the arrival rate is underestimated because the increase in surplus due to low prices is bounded by the number of available products.

Remark 5. *The revenue manager could also use variations in proposed prices to try to learn about the real value of the arrival rate during the booking process. She would then be able to update her pricing strategy using her estimated value for* $\tilde{\lambda}$*. This would increase the impact of revenue management on profits. Such a learning model falls beyond the scope of our paper.*

4.6 Heterogeneous Consumers and Products

We now turn to the case in which consumers are heterogeneous and have to choose between two types of products heterogeneous in their quality. Indeed, intertemporal price discrimination can seem unfair to consumers since two buyers can pay largely different prices for the same products. In particular, in the transport industry, business passengers are more likely to pay higher fares since they do not plan their trips in advance and buy their tickets when demand and prices are at their highest. For the same level of service, they might end up paying twice as much as a leisure passenger.

However, some papers in the literature argue that business passengers actually benefit from the presence of leisure passengers since the latter drive peak-demand prices down. (See Dana Jr [13].)

The situation we model here corresponds to the case in which a revenue manager optimizes profit over a single leg, i.e. same origin-destination, but she can offer two types of products, for instance tickets for a train early in the morning or for a train in the middle of the afternoon. In this case, high quality products are tickets for the more convenient train, which we call the rush-hour train because everyone wants to take it. The less convenient train is called the off-peak train.

In this section, we focus on a transport setting and test whether or not revenue management acts as a redistributive tool (and in which direction) in a general framework in which leisure passengers with a high price elasticity have a constant arrival rate through time and business passengers with a low price elasticity have an increasing arrival rate.[8]

Figure 4.6.1 depicts the unconditional probabilities of arrival of each type of consumers when demand is low. Our model for the dynamic arrival of consumers and the extension of the Bellman equation to heterogeneous consumers are thoroughly described in Appendix D.4.

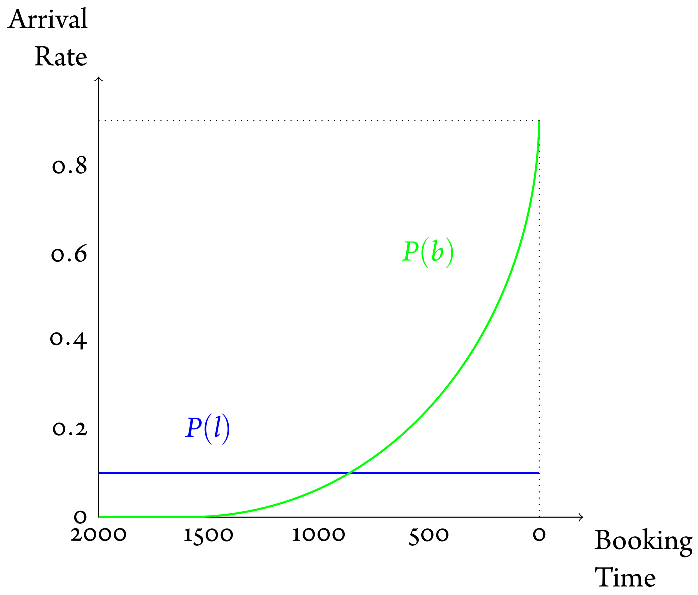

Figure 4.6.1: Change in the arrival rates of leisure and business passengers (unconditional probabilities) as a function of the booking period when the intensity of demand is low: $\mathbb{E}(\#arrival) = 400$.

Analysing revenue management as a redistribution tool requires to define different sets of parameters for each group of consumers. We consider here the case of business and leisure passengers. A summary of our simulation setup can be found in Table D.2.1 of Appendix D.2, column (2). The arrival rates are defined so as to induce an expected number of arrivals over the booking period:

$$\mathbb{E}(arrivals) \in \{400, 600, 800, 1000, 1200, 1400, 1600, 1800\}$$

In each case, the number of consumers' arrivals is in expectation equally divided between business and leisure passengers. Hence, leisure passengers are relatively more frequent at the beginning of the booking period and vice versa.

Both business and leisure passengers have identical preferences except for their price sensitivity. Using

[8] Business passengers can have a low price elasticity because their company covers their travel expenses for instance.

Equation 4.3.6, the price elasticity of demand is $\eta \approx -0.65$ for a business passenger facing a ticket for the more convenient train at price $p = 100$ and $\eta \approx -1.45$ for a leisure passenger facing a ticket for the less convenient train at price $p = 60$. As a reference, in their meta-analysis of price elasticities in U.K. rail transport industry, Wardman and Shires [56] find that the price elasticities of demand for inter-urban trains by London business passengers range from -0.54 to -0.63 in the short-run and from -0.79 to -0.92 in the long-run. The same elasticities for the London leisure passengers respectively range from -1 to -1.17 and from -1.47 to -1.72 in the short and long-run.

We compute the optimal fixed price vector $\mathbf{p^o}$ and the associated average profit and consumer surplus. We then compute the same average profit and consumer surplus in the case of a revenue management strategy using the algorithm presented in section 4.4. The choice set of the revenue manager is defined by:

$$P_i = \{30, 50, 80, 100, 120, 135, 150, 165, 180, 195, 210, 230, 250, 290\}\ ,\ i = 1, 2$$

This choice set covers a wide range of possible prices but is not optimized and has deliberately been chosen to be coarse.

The results are summarized in Tables 4.6.1 and 4.6.2. The impact of revenue management on profits is even more noticeable than in the previous sections as the effect is positive and significant whether the arrival of consumers is frequent or not.

Table 4.6.1 shows that revenue management significantly increases profits compared to an optimal fixed price strategy, even if the set of prices from which the revenue manager can choose has not been optimized. Indeed, revenue management can increase profits from 1.3% up to 13.7%, depending on the market size. The impact on profits is also larger for small market sizes, which greatly contrasts with our results on homogeneous consumers.

The variations in consumer surpluses between business and leisure passengers indicate that revenue management benefits leisure passengers at the expense of business ones for small market sizes. However, as the market size increases, the negative impact of revenue management on business surplus disappears and the effect even becomes significantly positive. The impact on leisure surplus on the other hand is reduced and becomes non-significant. For instance, when the expected number of arrivals is 1800, profits are increased by 1.3% due to revenue management, while the impact on business and leisure passengers is positive with their respective surplus increasing by 1.7% and 12%. The impact on leisure passengers is non-significant. When the expected number of arrival is 400, revenue management increases both profits and leisure consumer surplus by respectively 13.7% and 491%, and lowers business surplus by 14.5%. Revenue management increases social welfare compared to a fixed price strategy for sufficiently high intensities of demand ($\mathbb{E}(arrival) \geq 1600$) but has a more ambiguous impact for lower intensities as the

Table 4.6.1: Profit and consumer surplus under revenue management and a fixed price strategy for different market sizes and heterogeneous consumers.

$\mathbb{E}(arrivals)$	Pricing	$\mathbf{p^o}$	Avg. Profit	Avg. Surplus Business	Avg. Surplus Leisure
400	Fixed	135/135	16202	152	14
	RM		18421	130	86
	Gain from RM		13.7%**	−14.5%**	491%**
600	Fixed	135/135	24333	226	22
	RM		26841	207	99
	Gain from RM		10.3%**	−8.2%**	355%**
800	Fixed	135/135	32502	304	30
	RM		34763	286	95
	Gain from RM		7.0%**	−5.8%**	220%**
1000	Fixed	135/135	40624	378	38
	RM		42581	357	85
	Gain from RM		4.8%**	−5.5%**%	127%**
1200	Fixed	148/136	48186	422	34
	RM		49575	417	70
	Gain from RM		2.9%**	−1.1%**%	107%**
1400	Fixed	164/130	55655	472	35
	RM		56680	461	53
	Gain from RM		1.8%**	−2.3%**	51%**
1600	Fixed	182/140	62463	482	29
	RM		63449	479	31
	Gain from RM		1.6%**	−0.6%	6.5%
1800	Fixed	199/154	68826	479	20
	RM		69741	487	23
	Gain from RM		1.3%**	1.7%**	12%

Significance levels: † : 10%; ∗ : 5%; ∗∗ : 1%

increase in profit and leisure surplus comes at the expense of business passengers.

Table 4.6.2 gives us some intuition about why revenue management affects consumer surplus: for low values of demand, revenue management actually decreases the load rates for business passengers to the benefit of leisure passengers. Indeed, the shape of the arrival of the different types of consumers allows the revenue manager to discriminate easily between them, i.e. to fill the train with leisure passengers at the beginning of the booking period and to extract business surplus as much as possible when business demand is high. Higher values of demand call for discrimination between the two trains, as shown by the levels of the optimal fixed prices. The revenue manager wants to allocate leisure passengers to the off-peak train. Revenue management makes this discriminatory process more efficient, which leads to higher load rates, thus higher surplus.

Figures D.8.14 and D.8.15 of Appendix D.8.3 give us some additional evidence about the discriminatory process. First, Figure D.8.14 shows how often each price of the choice set P is chosen for each train. For each demand intensity, the revenue manager chooses several prices with positive probability through the booking period, which illustrates intertemporal price discrimination. However, for an expected number of arrivals inferior to 800, the distributions of chosen prices for the off-peak and rush-hour trains are almost identical: second-order price discrimination is suboptimal in that case. For an expected arrival greater than 1000, the distribution of prices for the rush-hour train slides to the right of the distribution for the off-peak train: differentiation between the off-peak and the rush-hour train becomes optimal when demand is sufficiently high, i.e. when resources (here sold seats) become limited.

Figure D.8.15 displays the dynamics of fares through the booking period and gives another evidence of this mechanism. For $\mathbb{E}(arrivals) = 600$, i.e. for a small market, the probability that there remains some empty seats in each train at the end of the booking period is almost one and prices evolve in an almost completely deterministic way. Intertemporal price discrimination is only driven by the increasing arrival of high-value consumers and posted prices are *de facto* independent of past sales. The increasing scale shape of the price function comes from the revenue manager trying to extract the business surplus as the arrival rate of business passengers increases.[9] For $\mathbb{E}(arrivals) = 1800$, the price curves are smooth and have the classic revenue management shape, indicating that, as the number of available seats is probably binding, posted prices are now non-monotonic functions of the previous sales.

[9]The small irregularities in the shape of the off-peak train price functions are due to the two opposite incentives of the revenue manager: as time goes by, business passengers are more frequent but filling the train also becomes more urgent.

Table 4.6.2: Comparison of load rates between revenue management and a fixed price strategy for different intensities of demand when consumers are heterogeneous.

$\mathbb{E}(arrivals)$	Pricing	Posted Prices	Avg. Load Rate Rush-Hour Train (B)	Avg. Load Rate Rush-Hour Train (L)	Avg. Load Rate Off-Peak Train (B)	Avg. Load Rate Off-Peak Train (L)
400	Fixed	135/135	33%	4%	20%	3%
	RM		29%	20%	18%	12%
	$\Delta(load)$		$-4\%^{**}$	$16\%^{**}$	$-2\%^{**}$	$10\%^{**}$
600	Fixed	135/135	50%	7%	30%	4%
	RM		46%	24%	28%	15%
	$\Delta(load)$		$-3\%^{**}$	$18\%^{**}$	$-2\%^{**}$	$11\%^{**}$
800	Fixed	135/135	66%	9%	40%	6%
	RM		63%	24%	38%	16%
	$\Delta(load)$		$-3\%^{**}$	$15\%^{**}$	$-2\%^{**}$	$10\%^{**}$
1000	Fixed	135/135	82%	11%	50%	7%
	RM		75%	21%	51%	16%
	$\Delta(load)$		$-7\%^{**}$	$10\%^{**}$	$0.8\%^{**}$	$9\%^{**}$
1200	Fixed	148/136	89%	9%	63%	8%
	RM		84%	14%	65%	18%
	$\Delta(load)$		$-5\%^{**}$	$5\%^{**}$	$2\%^{**}$	$11\%^{**}$
1400	Fixed	164/130	91%	7%	80%	11%
	RM		90%	8%	78%	17%
	$\Delta(load)$		$-0.5\%^{**}$	$1\%^{**}$	$-2\%^{**}$	$5\%^{**}$
1600	Fixed	182/140	93%	4%	87%	10%
	RM		94%	5%	86%	10%
	$\Delta(load)$		$1\%^{**}$	$0.3\%^{**}$	$-0.7\%^{**}$	$0.8\%^{**}$
1800	Fixed	199/154	95%	3%	90%	7%
	RM		96%	3%	91%	7%
	$\Delta(load)$		$1\%^{**}$	$0.2\%^{*}$	$0.5\%^{**}$	0.1%

Significance levels: † : 10%; $*$: 5%; $**$: 1%

4.7 Revenue Management and Intermodal Competition

In this section, we test how competition affects revenue management practices. As transport industry is the most likely economic sector featuring competition between revenue managers, we focus on the two following cases of competition:

- Indirect competition — for instance *Rail Vs. Road*: Larger road infrastructures reduce driving costs (either time or fuel consumption costs) and thus affect demand for trains. Since the building decisions concerning these infrastructures are mainly political and have long-term impacts on the market, strategic interactions between the revenue manager and the policy maker are non-existent and we model this setup using a reduced-form approach. The level of competition in that case is adjusted via the value of the outside option. We present this part of the analysis of competition in Appendix D.5.
- Direct competition — for instance *Rail Vs. Air*: As both railway and airline industries use revenue management and are substitutable for distances between 400 and 1000 kilometres, decisions of both revenue managers are highly strategic and must be studied through a game-theoretic approach. We use the concept of subgame-perfect Nash equilibrium to predict probable outcomes of these repeated strategic interactions.

4.7.1 The Revenue Management Game

Two revenue managers play a stochastic game with finite horizon. Each revenue manager proposes homogeneous products, which can be thought of as tickets for a transport mode, e.g. a plane or train ticket. However, we allow for heterogeneity between revenue managers. In the following, a type of products i is associated with a revenue manager. In that case, $-i$ denotes the products of the other revenue manager.

We assume that consumers are not forward-looking and maximize their utility when making their purchasing decision. We model demand using a multinomial logit approach. The utility of consumers is therefore defined along the lines of section 4.3.3:

$$\begin{cases} u^t(\theta^i) = \nu + a^i - \gamma p^i + \varepsilon^{it} & \text{, for product } i \text{ at price } p^i. \\ u^t_0 = \varepsilon^{0t} & \text{, when the outside option is chosen.} \end{cases} \tag{4.7.1}$$

We also assume that at any stage of the game the number of remaining units of each type of products is common knowledge among revenue managers.

The game is defined by:

- 2 revenue managers R and A, which respectively stand for railway and airline.

- The total capacities of both types of products: $\bar{X} = (\bar{X}^R, \bar{X}^A)'$.

- The booking period of length T.

- The state space $X_t \in T \times \{0, \dots, \bar{X}^R\} \times \{0, \dots, \bar{X}^A\}$, which represents any possible remaining capacity X at any date t of the booking process.

- At each period, each revenue manager proposes a price in $\tilde{P} = \{\underline{p}^i, \bar{p}^i, +\infty\}$ for $i = R, A$, where $\underline{p}^i$ and $\bar{p}^i$ respectively represent a low and a high price. When a revenue manager chooses $+\infty$, her product is bought with zero probability. It is direct that for any action and state space, action $\tilde{p}^i = +\infty$ for $i = R, A$ is dominated by posting a positive and finite price, unless $X^i = 0$. Indeed, posting the high price yields the maximum possible profit with a positive probability regardless of the competitor's action. Conversely, a revenue manager cannot choose $\tilde{p}^i$ finite when $X^i = 0$. The action $\tilde{p}^i$ chosen at each period by revenue manager i when $X^i \neq 0$ is therefore any element of $\tilde{P}$ or any probability distribution over $\{\underline{p}^i, \bar{p}^i\}$

- The transition probabilities between states $Pr_t : X_t \times \tilde{P} \longmapsto X_{t-1}$, defined by:

$$\begin{cases} Pr_t(X^i_{t-1} = X^i_t - 1, X^{-i}_{t-1} = X^{-i}_t | X_t, \tilde{p}) = \frac{e^{\bar{u}(\tilde{p}^i)}}{e^{\bar{u}(\tilde{p}^i)} + e^{\bar{u}(\tilde{p}^{-i})} + 1} \\ Pr_t(X_{t-1} = X_t | X_t, \tilde{p}) = \frac{1}{\sum_{i \in \{R,A\}} e^{\bar{u}(\tilde{p}^i)} + 1} \end{cases}, i = R, A;\ t = T, \dots, 1. \tag{4.7.2}$$

- The gain functions φ^i_t at date t for each revenue manager i. These gain functions correspond to the value functions in the monopolistic case. φ^i_t represents the profit of revenue manager i from date t onwards if strategies $(\tilde{p}_\tau)_{\tau \leq t}$ are played and remaining capacities are given by X:

$$\begin{aligned} \varphi^i_t(X, \tilde{p}_t) &= \lambda_t \left\{ Pr(X^i - 1, X^{-i} | X, \tilde{p}) \left[\tilde{p}^i + \varphi^i_{t-1}(X^i - 1, X^{-i}, \tilde{p}_{t-1}) \right] \right. \\ &\quad \left. + Pr(X^i, X^{-i} - 1 | X, \tilde{p}) \varphi^i_{t-1}(X^i, X^{-i} - 1, \tilde{p}_{t-1}) \right\} \\ &\quad + \left[1 - \lambda_t + \lambda_t Pr(X | X, \tilde{p}) \varphi^i_{t-1}(X, \tilde{p}_{t-1}) \right] \\ &= \lambda_t \left\{ Pr(X^i - 1, X^{-i} | X, \tilde{p}) \left[\tilde{p}^i + \varphi^i_{t-1}(X^i - 1, X^{-i}, \tilde{p}_{t-1}) - \varphi^i_{t-1}(X, \tilde{p}_{t-1}) \right] \right. \\ &\quad \left. + Pr(X^i, X^{-i} - 1 | X, \tilde{p}) \left[\varphi^i_{t-1}(X^i, X^{-i} - 1, \tilde{p}_{t-1}) - \varphi^i_{t-1}(X, \tilde{p}_{t-1}) \right] \right\} \\ &\quad + \varphi^i_{t-1}(X, \tilde{p}_{t-1}) \end{aligned} \tag{4.7.3}$$

Definition 7. *A subgame-perfect equilibrium of this game is defined by:*

- *Consumers maximize their utility when making their purchasing decision.*
- *At each period, each revenue manager plays her best response against the action of the other revenue manager.*

The existence of such an equilibrium is a standard result. This game can be solved by backward induction. Finding an subgame-perfect equilibrium of the game amounts to solving the following Bellman equation for $i = R, A$:

$$V_t^i(X, \tilde{p}^{-i}) = \max_{p^i} \varphi_t^{i*}(X, p^i, \tilde{p}^{-i}) \tag{4.7.4}$$

in which φ_t^{i*} is the gain function of player i at date t when a subgame-perfect equilibrium is played in the continuation game at $t - 1$.

To deal with the multiplicity of equilibria and properly solve this game, we restrict our set of equilibria *at each continuation game* in the following way:

- If only one equilibrium in pure strategies exists, it is played.
- If several equilibria in pure strategies exist, the one maximizing the joint payoff is played.
- If no equilibrium in pure strategies exists, both revenue managers use mixed strategies.[10]

4.7.2 Computation of the equilibrium

The only difference between the algorithms used for the duopoly and the monopoly lies in the way we compute the value function, as the action of one player is now influenced by the action of the other. We use backward induction to solve for the equilibrium of this finite sequential game. In the last period, the continuation values are zero and we compute the best response functions of all players.

If there exists a unique equilibrium in pure strategies, we can move up to the previous period using equilibrium profits as continuation values. If there are multiple equilibria, we select one according to the selection rule mentioned above. If only a mixed strategy equilibrium exists, we compute the mixing probabilities using the formula in Appendix D.6.

The generation of consumers arrival and purchasing decisions is straightforward except that we also draw a realization of the mixed strategies in the case of a mixed equilibrium.

[10]For a derivation of a mixed equilibrium of the game, see Appendix D.6.

4.7.3 Mono-Product Duopoly Vs. Multi-Product Monopoly

To investigate the impact of competition between revenue managers, we compare profits and consumer surpluses generated by a mono-product duopoly and a multi-product monopoly. By multi-product monopoly, we mean a situation in which one revenue manager controls the prices of both types of products. This allows us to isolate the impact of competition between revenue managers while preserving the structure of the market. As explained in more details in Appendix D.7, given our consumer preferences and multinomial logit approach, competition introduces a *de facto* horizontal differentiation between products, even when they have identical characteristics. Compared to a mono-product duopoly, a mono-product monopolist would therefore suppress this differentiation and reduce the total probability of purchasing a product relative to the outside option.

In the following, we assume that a homogeneous population of consumers is indifferent between the two types of products. Values of parameters used in this setup are summarized in Table D.2.1 of Appendix D.2, column (4). The sets of prices from which the duopoly and monopoly can choose are identical and equal to $\{40, 90\}$. Arrival rates belong to the following set: $\lambda \in \{0.2, 0.3, 0.4, 0.5, 0.7\}$.[11]

Note that the only differentiation here between the products is horizontal and the total capacity for each type is 200 units. Therefore, the profits generated by each revenue managers are perfectly equal. In Table 4.7.1, we only report the profit of one firm and not the joint profits of the industry. However, in addition to the consumer surplus generated by each firm, we choose to report the total consumer surplus generated during the booking period. For the multi-product monopoly, we report the monopoly profit and total consumer surplus.

Table 4.7.1 shows that, compared to the duopoly, the monopoly yields higher profits and reduces the consumer surplus through lower load rates. As shown in the following section, higher monopoly prices drive these results.

4.7.4 Price Dispersion and Competition

We study the impact of competition on price dispersion by comparing the average prices posted during the booking period and their average standard deviation for a multi-product monopoly and a mono-product duopoly. As we assume no vertical differentiation between products, results are perfectly symmetric.

Table 4.7.2 shows that competition reduces the average prices posted by revenue managers. In the

[11] For arrival rates greater than 0.7, both means of transportation are full with probability one by the end of the booking period and results are therefore similar to the case in which $\lambda = 0.7$.

Table 4.7.1: Profit, consumer surplus, and load rates for various intensities of demand and revenue management pricing with $P = \{40, 90\}$ in the mono-product duopoly and multi-product monopoly.

	Mono-product Duopoly				Multi-product Monopoly		
λ	Avg. Profit	Avg. Surplus	Avg. Tot. Surplus	Avg. Load Rate	Avg. Profit	Avg. Surplus	Avg. Load Rate
0.2	5829	191	522	73%	13516	188	38%
0.3	9021	236	679	99%	20360	281	57%
0.4	12929	166	589	99%	27063	375	75%
0.5	16638	111	523	99%	33649	475	94%
0.7	18000	91	497	100%	36000	497	100%

Note: The average profit, consumer surplus, and load rates for the duopoly correspond to the profit, consumer surplus and load rates generated by one firm.

Table 4.7.2: Comparison of average prices and their standard deviations between a multi-product monopoly and a mono-product duopoly for different intensities of demand.

	Mono-Product Duopoly		Multi-Product Monopoly	
λ	Avg. Price	Dispersion	Avg. Price	Dispersion
0.2	40	0	90	0
0.3	45.5	14.9	90	0
0.4	64.8	24.7	90	0
0.5	84.2	13.7	90	0
0.7	90	0	90	0

monopoly, price dispersion is low as the revenue manager always posts the high price. In the duopoly, when the arrival rate of consumers is intermediate, the revenue managers have incentives to switch between the low and high prices, and price dispersion is high.

For extreme values of demand, consumers should expect a low price dispersion as the revenue managers will only post either low prices or high ones.[12]

4.8 Strategic Decisions and Revenue Management

In this section, we explore to which extent revenue management can affect the strategic decisions of an operations manager. In particular, as revenue managers base their pricing decisions considering the number of unsold units, adjusting the total capacity can have an impact on the seller's profits and the consumer surplus. Hence, in a situation of competition between two revenue managers, such as the one exposed in section 4.7.1, is there such a thing as an optimal capacity?

4.8.1 The Capacity Game in the Mono-Product Duopoly

Before competing against each other, firms using revenue management pricing may be able to choose their capacity. In the transport industry, this might translate in choosing the size of planes that an airline company affects to a given origin-destination leg.

As in section 4.7.1, we suppose that revenue managers maximize their profit for a given total capacity. However, we consider an additional ex-ante stage during which two strategic decision makers choose their capacity constraint among a fixed set of choices $C = \{100, 150, 200, 250\}$. Solving for the equilibrium in the second stage of the game, i.e. the Revenue Management game, we get the payoff matrix and corresponding average consumer surpluses presented in Table 4.8.1 for the capacity game.

Unsurprisingly, the matrix is symmetric. The unique static subgame-perfect Nash equilibrium of the game is given by $C^* = (250, 250)$, i.e. both firms choose the highest possible capacity at equilibrium. This static equilibrium yields a total consumer surplus of 826.

However, this equilibrium is not Pareto optimal since both firms would benefit from a similar increase in their profit by playing $C = (150, 150)$, i.e. by restraining the capacity.[13] If firms were able to commit to limit their capacity to 150 seats, they would be able to secure an increase in profits of 7.3%. This increase

[12]Note that part of these results are due to the coarse choice sets we chose. They might differ if we include the optimal fixed price in the choice sets of the revenue managers.

[13]In fact, if we consider the repeated game, we can easily implement (150, 150) as an equilibrium, using the static Nash reversion, as long as decision makers are moderately patient:

$$\frac{13.1}{1-\delta} \geq 14.2 + \delta\frac{12.2}{1-\delta} \Leftrightarrow \delta \geq 0.55$$

Table 4.8.1: First-stage payoffs and consumer surplus in the static duopolistic capacity game.

	Payoffs ($\times 10^3$)					Consumer surplus			
	100	150	200	250		100	150	200	250
100	9/9	9/13.4	9/14.8	9/15.1	100	250	307	385	483
150	13.4/9	13.1/13.1	12.5/13.9	11.6/14.2	150	309	380	476	579
200	14.8/9	13.9/12.5	12.9/12.8	11.9/13.2	200	382	473	591	698
250	15.1/9	14.2/11.6	13.2/11.8	12.2/12.2	250	481	582	695	826

Note: The figures presented here are averaged across 500 simulations of the revenue management game with a price set $P = \{40, 90\}$ and $\lambda = 0.4$.

in profits is driven by a decrease in the importance of the time constraint relative to the capacity constraint. This leads to a lower opportunity cost of posting a high price and selling units more slowly.

4.8.2 Comparison with the Multi-Product Monopoly

We now compare this result to the case in which the revenue manager is in a situation of monopoly and prior to the booking period chooses the capacity constraints over her two types of products. To make the comparison clear, we choose the case of a multi-product monopoly and no vertical differentiation between the two types. The results are presented in Table 4.8.2.

Table 4.8.2: First stage payoffs and consumer surplus in the static monopolistic capacity game.

	Payoffs ($\times 10^3$)					Consumer surplus			
	100	150	200	250		100	150	200	250
100	18	22.4	23.8	24.1	100	251	309	386	480
150	22.4	26.3	26.8	26.8	150	310	370	398	403
200	23.8	26.8	27	27.1	200	385	398	378	378
250	24.1	26.8	27.1	27	250	482	405	375	378

Note: The figures presented here are averaged across 1000 simulations with a price set $P = \{40, 90\}$ and $\lambda = 0.4$.

Compared to the duopoly, the situation in which the capacities on both types of products are at their

In that case, the consumer surplus strongly decreases and is equal to 380. However, this story of tacit collusion falls beyond the scope of our paper.

maximum is no longer necessarily played. As long as at least 200 units of each type are available, profits are maximized.[14] Taking the example of transports and assuming that proposing larger trains involves larger fixed costs, the monopoly chooses to restrict the capacity.

We find an increase of 11% in overall profits compared to the duopoly. Consumer surplus on the other hand is less than half the consumer surplus when there is competition between revenue managers. Even in the case of collusion on the proposed capacity, the joint profit and consumer surplus would be similar to the ones under the monopoly, but for a reduced capacity. The welfare effects on competition between revenue managers are therefore positive.

4.9 Conclusion

Revenue management is the main pricing practice in the transport industry and as such, its impact on welfare deserves to be studied. Revenue management in itself has a number of properties which affect profits and consumer surplus compared to other pricing strategies. By offering more leeway to the seller in her choice set, it weakly increases profits compared to an optimal fixed price strategy. We also show in our computations that revenue management as an intertemporal price discrimination practice is only useful when there is a sufficiently high probability that profits are constrained by the number of available units. In this case, revenue management implies a higher consumer surplus compared to a situation in which only the optimal fixed price is posted. Such an increase comes from higher load rates and weakly lower average prices under revenue management. Coming back to transports, revenue management can increase welfare because more potential passengers can afford to travel. Revenue management is also a useful way to optimally respond to demand without actually having to optimize profits over the proposed set of prices. Even using a coarse set of prices to choose from and for different intensities of demand, the revenue manager succeeds in achieving at least as well as an optimal fixed price.

Applying our methodology to heterogeneous consumers, we are able to shed some light on the controversial issue of the *fairness* of revenue management practices. We find that for high values of demand, revenue management benefits all types of consumers. When demand is low, consumers with low price elasticities subsidize the ones with high price elasticities. Applied to transports, when demand is high, revenue management allows leisure passengers to buy cheap tickets without hurting business passengers. But, when demand is low, business passengers no longer benefit from the relatively lower prices they would pay under a fixed price policy.

In the second part of the paper, we study competition between revenue managers and find classic results: compared to a multi-product monopoly, revenue management in a mono-product duopoly lowers

[14] The differences in the last two rows and columns of Table 4.8.2 are not significant.

the average price and increases consumer surplus through higher load rates.

We also consider a setup in which revenue managers can choose their total capacity before competing against each other. We find that compared to the monopolistic case, competition between revenue managers does not reduce by much the joint profit of the industry but allows consumer surplus to be twice as high.

References

[1] Elizabeth S Allman, Catherine Matias, and John A Rhodes. Identifiability of parameters in latent structure models with many observed variables. *The Annals of Statistics*, 37(6A):3099–3132, 2009.

[2] Eric T Anderson, Nanda Kumar, and Surendra Rajiv. A comment on:"revisiting dynamic duopoly with consumer switching costs". *Journal of Economic Theory*, 116(1):177–186, 2004.

[3] Suman Basuroy and S Abraham Avri Ravid. How relevant are experts in the internet age? evidence from the motion pictures industry? 2013.

[4] Suman Basuroy, Subimal Chatterjee, and S Abraham Ravid. How critical are critical reviews? the box office effects of film critics, star power, and budgets. *Journal of Marketing*, pages 103–117, 2003.

[5] Alan Beggs and Paul Klemperer. Multi-period competition with switching costs. *Econometrica: Journal of the Econometric Society*, pages 651–666, 1992.

[6] Moshe E Ben-Akiva and Steven R Lerman. *Discrete choice analysis: theory and application to travel demand*, volume 9. MIT press, 1985.

[7] G. Biglaiser and J. Crémer. Equilibria in an infinite horizon game with an incumbent, entry and switching costs. *International Journal of Economic Theory*, 7(1):65–75, 2011.

[8] Gary Biglaiser, Jacques Crémer, and Gergely Dobos. The value of switching costs. *Journal of Economic Theory*, 148(3):935–952, 2013.

[9] Peter Boatwright, Suman Basuroy, and Wagner Kamakura. Reviewing the reviewers: The impact of individual film critics on box office performance. *Quantitative Marketing and Economics*, 5(4): 401–425, 2007.

[10] Luis Cabral. Switching costs and equilibrium prices. 2012.

[11] F. Camara and N. Dupuis. Avoiding judgement by recommending inaction: Beliefs manipulation and reputational concerns. 2013.

[12] Ramon Caminal and Carmen Matutes. Endogenous switching costs in a duopoly model. *International Journal of Industrial Organization*, 8(3):353–373, 1990.

[13] J.D. Dana Jr. Using yield management to shift demand when the peak time is unknown. *The Rand Journal of Economics*, pages 456–474, 1999.

[14] P. Dasgupta and E. Maskin. The existence of equilibrium in discontinuous economic games, i: Theory. *The Review of Economic Studies*, 53(1):1–26, 1986.

[15] S. DellaVigna and J. Hermle. Does media concentration lead to biased coverage? evidence from movie reviews. 2014.

[16] Loretti Dobrescu, Michael Luca, and Alberto Motta. What makes a critic tick? connected authors and the determinants of book reviews. 2012.

[17] Nicolas Eber. Switching costs and implicit contracts. *Journal of Economics*, 69(2):159–171, 1999.

[18] L. Einav. Seasonality in the us motion picture industry. *The Rand journal of economics*, 38(1): 127–145, 2007.

[19] Natalia Fabra and Alfredo García. Dynamic price competition with switching costs. 2012.

[20] J. Farrell and C. Shapiro. Dynamic competition with switching costs. *The RAND Journal of Economics*, pages 123–137, 1988.

[21] Joseph Farrell. A note on inertia in market share. *Economics Letters*, 21(1):73–75, 1986.

[22] Guillermo Gallego and Garrett Van Ryzin. A multiproduct dynamic pricing problem and its applications to network yield management. *Operations Research*, 45(1):24–41, 1997.

[23] M. Gentzkow and J.M. Shapiro. Media bias and reputation. *Journal of Political Economy*, 114(2): 280–316, April 2006.

[24] A. Gershkov and B. Moldovanu. Dynamic revenue maximization with heterogeneous objects: A mechanism design approach. *American Economic Journal: Microeconomics*, 1(2):168–198, 2009.

[25] Phil B Goodwin. A review of new demand elasticities with special reference to short and long run effects of price changes. *Journal of transport economics and policy*, pages 155–169, 1992.

[26] Bengt Holmström. Managerial incentives problems: A dynamic perspective. *The Review of Economic Studies*, 66(1):169–182, 1999.

[27] J. Hörner and L. Samuelson. Managing strategic buyers. *The Journal of Political Economy*, 119(3): 379–425, 2011.

[28] Matias Iaryczower and Matthew Shum. The value of information in the court: Get it right, keep it tight. *The American Economic Review*, 102(1):202–237, 2012.

[29] P. Klemperer. Markets with consumer switching costs. *The Quarterly Journal of Economics*, 102(2): 375–394, 1987.

[30] Paul Klemperer. The competitiveness of markets with switching costs. *The RAND Journal of Economics*, pages 138–150, 1987.

[31] Paul Klemperer. Price wars caused by switching costs. *The Review of Economic Studies*, 56(3): 405–420, 1989.

[32] John Lazarev. The welfare effects of intertemporal price discrimination: an empirical analysis of airline pricing in us monopoly markets. *New York University, working paper*, 2012.

[33] Tracy R Lewis and Huseyin Yildirim. Managing switching costs in multiperiod procurements with strategic buyers. *International Economic Review*, 46(4):1233–1269, 2005.

[34] Beatriz Mariano. Market power and reputational concerns in the ratings industry. *Journal of Banking & Finance*, 2012.

[35] R.P. McAfee and V. Velde. Dynamic pricing with constant demand elasticity. *Production and Operations Management*, 17(4):432–438, 2008.

[36] J.I. McGill and G.J. Van Ryzin. Revenue management: Research overview and prospects. *Transportation science*, 33(2):233, 1999.

[37] Bolandifar E. Modarres M. Dynamic pricing with periodic review and a finite set of prices with cancellation. *Journal of Industrial and Systems Engineering (JISE)*, 2009.

[38] S. Morris. Political correctness. *Journal of Political Economy*, 109(2):231–265, April 2001. URL http://ideas.repec.org/a/ucp/jpolec/v109y2001i2p231-265.html.

[39] Chakravarthi Narasimhan. Competitive promotional strategies. *Journal of Business*, pages 427–449, 1988.

[40] Tore Nilssen. Two kinds of consumer switching costs. *The RAND Journal of Economics*, pages 579–589, 1992.

[41] M. Ottaviani and P.N. Sørensen. Information aggregation in debate: who should speak first? *Journal of Public Economics*, 81(3):393–421, 2001.

[42] M. Ottaviani and P.N. Sørensen. Reputational cheap talk. *The Rand journal of economics*, 37(1): 155–175, 2006.

[43] A Jorge Padilla. Mixed pricing in oligopoly with consumer switching costs. *International Journal of Industrial Organization*, 10(3):393–411, 1992.

[44] A.J. Padilla. Revisiting dynamic duopoly with consumer switching costs. *Journal of Economic Theory*, 67(2):520–530, 1995.

[45] D.A. Reinstein and C.M. Snyder. The influence of expert reviews on consumer demand for experience goods: A case study of movie critics. *The journal of industrial economics*, 53(1):27–51, 2005. ISSN 1467-6451.

[46] Andrew Rhodes. Re-examining the effects of switching costs. *Economic Theory*, 57(1):161–194, 2014.

[47] Robert W Rosenthal. A dynamic model of duopoly with customer loyalties. *Journal of Economic Theory*, 27(1):69–76, 1982.

[48] Robert W Rosenthal. Dynamic duopoly with incomplete customer loyalties. *International Economic Review*, pages 399–406, 1986.

[49] Y. Shilony. Mixed pricing in oligopoly. *Journal of Economic Theory*, 14(2):373–388, 1977.

[50] N.L. Stokey. Intertemporal price discrimination. *The Quarterly Journal of Economics*, pages 355–371, 1979.

[51] K. Talluri and G. Van Ryzin. Revenue management under a general discrete choice model of consumer behavior. *Management Science*, pages 15–33, 2004.

[52] K.T. Talluri and G. Van Ryzin. *The theory and practice of revenue management*, volume 68. Springer Verlag, 2005.

[53] Theodore To. Multi-period competition with switching costs: An overlapping generations formulation. *The Journal of Industrial Economics*, pages 81–87, 1996.

[54] J Miguel Villas-Boas. Notes on switching costs and dynamic competition. *UC Berkeley*, 2011.

[55] G. Vulcano, G. van Ryzin, and W. Chaar. Om practice—choice-based revenue management: An empirical study of estimation and optimization. *Manufacturing & Service Operations Management*, 12(3):371–392, 2010.

[56] Mark Wardman and JD Shires. Review of fares elasticities in great britain. 2003.

[57] Kevin R Williams. Dynamic airline pricing and seat availability. Technical report, mimeo, 2013.

A
Appendix — Chapter 1

A.1 Micro-foundations for the Precision of the Aggregate Signal τ

Let us focus on the case of an experience good. The state of the world here is the quality of the product. Suppose the expert sends her recommendation to N potential consumers or agents. Each of the N agents is defined by an idiosyncratic taste $\varepsilon \sim U[0,1]$. They observe the recommendation of the expert and their own taste to decide whether or not to purchase the product.

Agents derives their utility from their own experience when they consume the product. As the expert, this experience is modelled as a private noisy signal. The precision of the agents' signal is denoted $l^a \in \left(\frac{1}{2}, 1\right]$. As this precision is greater than $\frac{1}{2}$, agents have a higher chance to get a good experience when they consume a high quality product than when they consume a low quality product. Ex-ante, they prefer to consume products with the higher expected quality. A potential consumer decides to purchase if her expected utility, is greater than the price normalized to 1:

$$\mathbb{E}(u) = l^a p(\theta_1|message) + (1-l^a)p(\theta_0|message) + \varepsilon = l^a(2p(\theta_1|message) - 1) + 1 + \varepsilon > 1$$

Suppose that n agents out of the N potential consumers purchase the product. Clearly $\mathbb{E}(n)$ increases with the prior on the high state, $p(\theta_1|message)$.

The n agents then individually experience the product. They have an experience corresponding to the true state of the world with probability l^a. We assume that all agents in the pool share their experience with the others. If n is large, it is easy to get a good idea about the true quality. If the agents have an experience corresponding to the true quality with probability 0.7 and out of 1000 agents who purchased the product, 700 had a good experience, one can be fairly certain than the quality was high. But if only 1 agent had purchased the same product, this agent would have had a 30% probability to get a bad experience. Her only report would not have permit to be definite about the quality of the product.

Overall, in this example, a higher updated prior on the high state of the world implies a greater number of agents who observe a signal on this state of the world, and in turn create a more precise aggregate signal, which justify our assumption that $\tau(v_1)$ is increasing.

A.2 Proof of Proposition 2

Let us first rewrite the characterization of the incentive compatible set. To simplify notations, we omit t and $\tilde{t}$ in $p^e(\tilde{\theta}_i|s_i, \tilde{s}_i, t, \tilde{t})$:

$$p^e(\tilde{\theta}_i|s_i, \tilde{s}_i) \geq p^e(\tilde{\theta}_{-i}|s_i, \tilde{s}_{-i})$$

$$\Leftrightarrow p(X_i|s_i, \tilde{s}_i)p(\tilde{\theta}_i|X_i, \tilde{s}_i) + p(X_{-i}|s_i, \tilde{s}_i)p(\tilde{\theta}_i|X_{-i}, \tilde{s}_i) \geq$$
$$p(X_i|s_i, \tilde{s}_{-i})p(\tilde{\theta}_{-i}|X_i, \tilde{s}_{-i}) + p(X_{-i}|s_i, \tilde{s}_{-i})p(\tilde{\theta}_{-i}|X_{-i}, \tilde{s}_{-i})$$

$$\Leftrightarrow \left[(2\tau(v(\tilde{s}_i)) - 1)p(\theta_i|s_i, t) + 1 - \tau(v(\tilde{s}_i))\right] p(\tilde{\theta}_i|X_i, \tilde{s}_i) +$$
$$\left[(1 - 2\tau(v(\tilde{s}_i)))p(\theta_i|s_i, t) + \tau(v(\tilde{s}_i))\right] p(\tilde{\theta}_i|X_{-i}, \tilde{s}_i) \geq$$
$$\left[(2\tau(v(\tilde{s}_{-i})) - 1)p(\theta_i|s_i, t) + 1 - \tau(v(\tilde{s}_{-i}))\right] p(\tilde{\theta}_{-i}|X_i, \tilde{s}_{-i}) +$$
$$\left[(1 - 2\tau(v(\tilde{s}_{-i})))p(\theta_i|s_i, t) + \tau(v(\tilde{s}_{-i}))\right] p(\tilde{\theta}_{-i}|X_{-i}, \tilde{s}_{-i})$$

$$\Leftrightarrow p(\theta_i|s_i, t) \left\{\left[2\tau(v(\tilde{s}_i)) - 1\right] \left(p(\tilde{\theta}_i|X_i, \tilde{s}_i) - p(\tilde{\theta}_i|X_{-i}, \tilde{s}_i)\right) +\right.$$
$$\left.\left[2\tau(v(\tilde{s}_{-i})) - 1\right] \left(p(\tilde{\theta}_{-i}|X_{-i}, \tilde{s}_{-i}) - p(\tilde{\theta}_{-i}|X_i, \tilde{s}_{-i})\right)\right\} \geq$$
$$(1 - \tau(v(\tilde{s}_{-i})))p(\tilde{\theta}_{-i}|X_i, \tilde{s}_{-i}) + \tau(v(\tilde{s}_{-i}))p(\tilde{\theta}_{-i}|X_{-i}, \tilde{s}_{-i}) -$$
$$(1 - \tau(v(\tilde{s}_i)))p(\tilde{\theta}_i|X_i, \tilde{s}_i) - \tau(v(\tilde{s}_i))p(\tilde{\theta}_i|X_{-i}, \tilde{s}_i)$$

The right-hand side does not depend on the true ability t of the expert and the expression between brackets in the left-hand side is positive and does not depend on the true ability either. Hence, the inequality is more slack when $p(\theta_i|s_i, t)$ increases. Therefore, if the incentive compatible condition is satisfied for an ability $l < h$, it is also satisfied for h.

A.3 Proof of Proposition 3

The proof for $\mu_1 \in IC_{l,\tilde{t}}$ is direct for both types. For $\mu \in [0, 1] \setminus IC_{l,\tilde{t}}$, we first prove the result for $\mu_1 > \bar{\mu}_1$ when $t = l$, then move to $t = h$. The result's proof when $\mu < \underline{\mu}_1$ is similar. This proof is a generalization of Ottaviani and Sørensen [41].

Let us introduce some notations which will simplify the demonstration. We denote $\breve{\sigma}$ the strategy described in the Proposition 3. According to $\breve{\sigma}$, for $\mu_1 > \bar{\mu}_1$, the low type expert truthfully reveals s_1 and randomizes after receiving s_0. We denote $\xi = p(m_0|s_0)$. We want to prove ξ exists and belongs to $(0, 1)$. If the low type expert randomizes after receiving s_0, it must be that:

$$p(h|m_0, \breve{\sigma}, s_0, l) = p(h|m_1, \breve{\sigma}, s_0, l) \tag{A.3.1}$$

where:

$$p(h|m, \breve{\sigma}, s_0, l) = \sum_{i \in \{0,1\}} p(X_i|s_0, l) \left(\sum_{j \in \{0,1\}} p(h)_{m,j}\, p(\theta_j|X_i, m, \breve{\sigma}) \right)$$

In the previous expression:

$$p(h)_{m,j} = p(h|m, \breve{\sigma}, \theta_j) = \begin{cases} \frac{hp(h)}{hp(h)+\xi lp(l)} & \text{if } (m,j) = (0,0) \\ \\ \frac{(1-h)p(h)}{(1-h)p(h)+\xi(1-l)p(l)} & \text{if } (m,j) = (0,1) \\ \\ \frac{hp(h)}{hp(h)+[l+(1-l)(1-\xi)]p(l)} & \text{if } (m,j) = (1,1) \\ \\ \frac{(1-h)p(h)}{(1-h)p(h)+[(1-l)+l(1-\xi)]p(l)} & \text{if } (m,j) = (1,0) \end{cases} \tag{A.3.2}$$

Since in equilibrium, the agents correctly forecast the expert's strategy, Equation A.3.1 can be rewritten:

$$\begin{aligned} & p(h)_{1,1} p^e(\theta_1^a|s_0, l, m_1, \breve{\sigma}) + p(h)_{1,0}(1 - p^e(\theta_1^a|s_0, l, m_1, \breve{\sigma})) \\ =& p(h)_{0,0} p^e(\theta_0^a|s_0, l, m_0, \breve{\sigma}) + p(h)_{0,1}(1 - p^e(\theta_0^a|s_0, l, m_0, \breve{\sigma})) \end{aligned} \tag{A.3.3}$$

- When $\xi = 1$, i.e. when the low type is always truthtelling, using Equations A.3.2, we get: $p(h)_{1,1} = p(h)_{0,0} > p(h)_{0,1} = p(h)_{1,0}$. Also, $p^a(\theta|X, m, \breve{\sigma}) = p^a(\theta|X, m, \sigma^T)$, hence:

  $$p^e(\theta^a|s, l, m, \breve{\sigma}) = p(X_1|s, l) p^a(\theta|X_1, \tilde{s}) + p(X_0|s, l) p^a(\theta|X_0, \tilde{s})$$

  where, as before, $\tilde{s} = m \cap \sigma^T$. Since, $\mu_1 \notin IC_{l,\tau}$ and $\mu_1 > \bar{\mu}_1$, necessarily:

  $$\begin{aligned} & p(X_1|s_0, l)\, [p(\theta_1|X_1, \tilde{s}_1) - p(\theta_0|X_1, \tilde{s}_0)] > \\ & p(X_0|s_0, l)\, [p(\theta_0|X_0, \tilde{s}_0) - p(\theta_1|X_0, \tilde{s}_1)] \\ \Leftrightarrow & p^e(\theta_1|s_0, l, m_1, \breve{\sigma}) > p^e(\theta_0|s_0, l, m_0, \breve{\sigma}) \end{aligned}$$

  Therefore $p(h|m_0, \breve{\sigma}, s_0, l) < p(h|m_1, \breve{\sigma}, s_0, l)$ for $\xi = 1$.

- When $\xi = 0$ (i.e. the low type always sends m_1), $p(h)_{0,1} = p(h)_{0,0} = 1$ so $p(h|m_0, \breve{\sigma}, s_0, l) = 1$ and $p(h)_{1,1}, p(h)_{1,0} < 1$ (unless $p(h) = 1$), therefore $p(h|m_1, \breve{\sigma}, s_0, l) < 1$.

Since $p(h|m_0, \breve{\sigma}, s_0, l) < p(h|m_1, \breve{\sigma}, s_0, l)$ and $p(h|m_0, \breve{\sigma}, s_0, l) > p(h|m_1, \breve{\sigma}, s_0, l)$ are not consistent at equilibrium with respectively $\xi = 1$ and $\xi = 0$, by continuity of $p(h|m, \breve{\sigma}, s_0, l)$ in ξ, there must exist $\xi \in (0,1)$ such that $p(h|m_0, \breve{\sigma}, s_0, l) = p(h|m_1, \breve{\sigma}, s_0, l)$.

We now need to check that it is optimal for the high type to follow $\breve{\sigma}$, i.e. to be truthtelling. Using Equation A.3.3, we have:

$$\begin{aligned}&(p(h)_{1,1} - p(h)_{1,0})p^e(\theta_1^a|s_0, l, m_1, \breve{\sigma}) + p(h)_{1,0}\\=&(p(h)_{0,0} - p(h)_{0,1})p^e(\theta_0^a|s_0, l, m_0, \breve{\sigma}) + p(h)_{0,1}\end{aligned}$$

Since $p(h)_{1,1} > p(h)_{1,0}$ and $p(h)_{0,0} > p(h)_{0,1}$, to prove the result, we only need to show:

$$\begin{cases} p^e(\theta_1^a|s_0, h, m_1, \breve{\sigma}) < p^e(\theta_1^a|s_0, l, m_1, \breve{\sigma}) \\ p^e(\theta_0^a|s_0, h, m_0, \breve{\sigma}) > p^e(\theta_0^a|s_0, l, m_0, \breve{\sigma}) \end{cases}$$

Checking it is easy as $p(\theta_0|s_0, h) > p(\theta_0|s_0, l)$.

A.4 Generalization of the results

Lemma 5 illustrates how the relative impact of the expert's message on the final posterior decreases with the precision of the aggregate signal.

Lemma 5. *For all $i, \tilde{s}, s, \tilde{t}, \hat{t}$, and for all functions $\tau : [0, 1] \to [1/2, 1]$, there exists a function $\alpha^{\tilde{s}}(.)$ such that:*

$$p^e(\theta_i^c|\tilde{s}, s, \tilde{t}, \hat{t}) = (1 - \alpha^{\tilde{s}}(\tau))\nu_i(\tilde{s}, \tilde{t}) + \alpha^{\tilde{s}}(\tau)p^e(\theta_i|s, \hat{t})$$

Moreover, $\alpha^{\tilde{s}}(1/2) = 0,\ \alpha^{\tilde{s}}(1) = 1,$ and $\frac{\partial \alpha^{\tilde{s}}}{\partial \tau} \geq 0$.

Proof. We have to show that $p^e(\theta_i^a|\tilde{s}, s) = (1 - \alpha(\tau))p^a(\theta_i|\tilde{s}) + \alpha(\tau)p^e(\theta_i|s)$ with α increasing in τ and $\alpha(1/2) = 0, \alpha(1) = 1$.

First, let us rewrite $p^e(\theta_i^a|\tilde{s}, s)$ as a weighted average of $p^a(\theta_i|\tilde{s})$ and $p^e(\theta_i|s)$. To save space we replace $\tau(\nu_1(\tilde{s}, \tilde{t}))$ by τ:

$$
\begin{aligned}
p^e(\theta_i^a|\tilde{s}, s) =& p^e(X_i|\tilde{s}, s)p^a(\theta_i|\tilde{s}, X_i) + p^e(X_j|\tilde{s}, s)p^a(\theta_i|\tilde{s}, X_{-i}) \\
=& \left[\tau p^e(\theta_i|s) + (1-\tau)p^e(\theta_{-i}|s)\right] p^a(\theta_i|\tilde{s}, X_i) \\
& + \left[\tau p^e(\theta_{-i}|s) + (1-\tau)p^e(\theta_i|s)\right] p^a(\theta_i|\tilde{s}, X_{-i}) \\
=& p^e(\theta_i|s)(2\tau - 1)\left[p^a(\theta_i|\tilde{s}, X_i) - p^a(\theta_i|\tilde{s}, X_{-i})\right] \\
& + (1-\tau)p^a(\theta_i|\tilde{s}, X_i) + \tau p^a(\theta_i|\tilde{s}, X_{-i}) \\
=& p^e(\theta_i|s)(2\tau - 1)\left[\frac{\tau p^a(\theta_i|\tilde{s})}{p^a(X_i|\tilde{s})} - \frac{(1-\tau)p^a(\theta_i|\tilde{s})}{p^a(X_{-i}|\tilde{s})}\right] \\
& + p^a(\theta_i|\tilde{s})(1-\tau)\tau\left[\frac{1}{p^a(X_i|\tilde{s})} + \frac{1}{p^a(X_{-i}|\tilde{s})}\right] \\
=& p^e(\theta_i|s)(2\tau - 1)\frac{p^a(\theta_i|\tilde{s})p^a(\theta_{-i}|\tilde{s})\left[\tau^2 - (1-\tau)^2\right]}{p^a(X_i|\tilde{s})p^a(X_{-i}|\tilde{s})} \\
& + p^a(\theta_i|\tilde{s})\tau(1-\tau)\left[\frac{p^a(X_i|\tilde{s}) + p^a(X_{-i}|\tilde{s})}{p^a(X_i|\tilde{s})p^a(X_{-i}|\tilde{s})}\right] \\
=& p^e(\theta_i|s)(2\tau - 1)^2\frac{p^a(\theta_i|\tilde{s})p^a(\theta_{-i}|\tilde{s})}{p^a(X_i|\tilde{s})p^a(X_{-i}|\tilde{s})} + p^a(\theta_i|\tilde{s})\frac{\tau(1-\tau)}{p^a(X_i|\tilde{s})p^a(X_{-i}|\tilde{s})}
\end{aligned}
$$

where $p^a(X_i|\tilde{s}) = \tau p(\theta_i|\tilde{s}) + (1-\tau)p(\theta_{-i}|\tilde{s})$.
We now pose $\alpha^{\tilde{s}}(\tau) = (2\tau-1)^2 \frac{p^a(\theta_i|\tilde{s})p^a(\theta_{-i}|\tilde{s})}{p^a(X_i|\tilde{s})p^a(X_{-i}|\tilde{s})}$.
To complete the proof, we need to show:

(a) $1 - \alpha^{\tilde{s}}(\tau) = \frac{\tau(1-\tau)}{p^a(X_i|\tilde{s})p^a(X_{-i}|\tilde{s})}$

(b) $\alpha^{\tilde{s}}(1/2) = 0$ and $\alpha^{\tilde{s}}(1) = 1$

(c) $\frac{\partial \alpha^{\tilde{s}}}{\partial \tau} \geq 0$

(a) is equivalent to proving:

$$
(2\tau-1)^2 p^a(\theta_i|\tilde{s})p^a(\theta_{-i}|\tilde{s}) = p^a(X_i|\tilde{s})p^a(X_{-i}|\tilde{s}) - \tau(1-\tau)
$$

which is indeed the case:

$$\begin{aligned}
&p^a(X_i|\tilde{s})p^a(X_{-i}|\tilde{s}) - \tau(1-\tau) \\
=&\tau(1-\tau)\left[p^a(\theta_i|\tilde{s})^2 + p^a(\theta_{-i}|\tilde{s})^2 - 1\right] + p^a(\theta_i|\tilde{s})p^a(\theta_{-i}|\tilde{s})\left[\tau^2 + (1-\tau)^2\right] \\
=&\tau(1-\tau)\left[2p^a(\theta_i|\tilde{s})^2 + 1 - 2p^a(\theta_i|\tilde{s}) - 1\right] + p^a(\theta_i|\tilde{s})p^a(\theta_{-i}|\tilde{s})\left[\tau^2 + (1-\tau)^2\right] \\
=&\tau(1-\tau)\left[-2p^a(\theta_i|\tilde{s})p^a(\theta_{-i}|\tilde{s})\right] + p^a(\theta_i|\tilde{s})p^a(\theta_{-i}|\tilde{s})\left[\tau^2 + (1-\tau)^2\right] \\
=&p^a(\theta_i|\tilde{s})p^a(\theta_{-i}|\tilde{s})\left[\tau - (1-\tau)\right]^2
\end{aligned}$$

(b) Direct as $\alpha^{\tilde{s}}(1/2) = 0 = (1-\alpha^{\tilde{s}})(1)$.
(c)

$$\begin{aligned}
&\frac{\partial \alpha^{\tilde{s}}}{\partial \tau} \geq 0 \\
\Leftrightarrow &4(2\tau-1)p^a(X_i|\tilde{s})p^a(X_{-i}|\tilde{s}) - (2\tau-1)^2 \frac{\partial p^a(X_i|\tilde{s})p^a(X_{-i}|\tilde{s})}{\partial \tau} \geq 0 \\
\Leftrightarrow &4p^a(X_i|\tilde{s})p^a(X_{-i}|\tilde{s}) \\
&-(2\tau-1)(2p^a(\theta_i|\tilde{s})-1)\left[(2\tau-1)(1-p^a(\theta_i|\tilde{s})) + (1-2\tau)p^a(\theta_i|\tilde{s})\right] \geq 0 \\
\Leftrightarrow &4p^a(X_i|\tilde{s})p^a(X_{-i}|\tilde{s}) + (2\tau-1)^2(2p^a(\theta_i|\tilde{s})-1)^2 \geq 0
\end{aligned}$$

□

The incentive compatibility condition for s_i is then equivalent to:

$$(1-\alpha^{\tilde{s}_i}(\tau_{\tilde{s}_i}))v_i(\tilde{s}_i) + \alpha^{\tilde{s}_i}(\tau_{\tilde{s}_i})p^e(\theta_i|s_i) \geq (1-\alpha^{\tilde{s}_{-i}}(\tau_{\tilde{s}_{-i}}))v_{-i}(\tilde{s}_{-i}) + \alpha^{\tilde{s}_{-i}}(\tau_{\tilde{s}_{-i}})p^e(\theta_{-i}|s_i) \qquad \text{(A.4.1)}$$

A.4.1 Extension of the Results with Unknown Ability

In this section, we focus on the case in which the expert and agents share the same information about the ability of the expert, i.e.:

$$\hat{t} = \tilde{t} = hp(h) + lp(l)$$

In that case, the asymmetry of information consists solely on the private signal received by the expert. Hence, if the expert decides to truthfully reveal this private signal, the asymmetry of information is lifted and the private posterior of the expert is equal to the updated prior of the agents:

$$v_i(\tilde{s}_i, \tilde{t}) = p^e(\theta_i|s_i, \tilde{t}) \qquad \text{(A.4.2)}$$

The incentive compatibility condition given by equation A.4.1 now simplifies to:

$$p^e(\theta_i|s_i) \geq (1-\alpha^{\tilde{s}_{-i}}(\tau_{\tilde{s}_{-i}}))v_{-i}(\tilde{s}_{-i}) + \alpha^{\tilde{s}_{-i}}(\tau_{\tilde{s}_{-i}})p^e(\theta_{-i}|s_i) \qquad \text{(A.4.3)}$$

The first thing that we can point out is that, in this setting, it is never optimal for the expert to lie and announce m_{-i} when she has received s_i and $p(\theta_i) \geq 1/2$. Indeed, $p(\theta_i) \geq 1/2 \Rightarrow p^e(\theta_i|s_i) \geq p^e(\theta_{-i}|s_i)$ and a sufficient condition for equation A.4.3 to be satisfied is $p^e(\theta_i|s_i) \geq v_{-i}(\tilde{s}_{-i})$, which is also satisfied when $p(\theta_i) \geq 1/2$. In other terms, when the expert receives a message confirming her expectations about the state of the world, she is better off by telling the truth. This also means that the prior $\mu_i = 1/2$ always belongs to the incentive compatibility set.

When the expert does not know her type, the truthful revelation set is equal to the incentive compatible set.

Proof of Proposition 5

Proof. It suffices to notice that the left-hand side of equation A.4.3 is invariant with τ, but that the right-hand side when using the function τ_1 is greater than when using the function τ_2. Indeed, since $\alpha(.)^{\tilde{s}_{-i}}$ is increasing:
$p(\theta_{-i}|\tilde{s}_{-i}) > p(\theta_{-i}|s_i) \Rightarrow a^{\tilde{s}_{-i}}(\tau_1(x)) \left[p(\theta_{-i}|s_i) - p(\theta_{-i}|\tilde{s}_{-i})\right] > a^{\tilde{s}_{-i}}(\tau_2(x)) \left[p(\theta_{-i}|s_i) - p(\theta_{-i}|\tilde{s}_{-i})\right]$.
Hence, by transitivity, if $\mu_1 \in IC_{\tilde{t},\tau_1}$, then $\mu_1 \in IC_{\tilde{t},\tau_2}$. □

To show the result on the barrier to entry effect, we first prove the following Lemma:

Lemma 6. $a^1(\tau) > a^0(\tau) \Leftrightarrow \mu_1 < \frac{1}{2}$

Proof. Recall that:

$$a^{\tilde{s}}(\tau) = (2\tau - 1)^2 \frac{p(\theta_i|\tilde{s})p(\theta_{-i}|\tilde{s})}{p(X_i|\tilde{s})p(X_{-i}|\tilde{s})}$$

Hence:

$$\begin{aligned}
&a^1(\tau) - a^0(\tau) > 0 \\
\Leftrightarrow &p(\theta_i|\tilde{s}_1)p(\theta_{-i}|\tilde{s}_1)p(X_i|\tilde{s}_0)p(X_{-i}|\tilde{s}_0) > p(\theta_i|\tilde{s}_0)p(\theta_{-i}|\tilde{s}_0)p(X_i|\tilde{s}_1)p(X_{-i}|\tilde{s}_1) \\
\Leftrightarrow &p(\theta_i|\tilde{s}_1)p(\theta_{-i}|\tilde{s}_1) \left(p(\theta_i|\tilde{s}_0)^2 + p(\theta_{-i}|\tilde{s}_0)^2\right) > p(\theta_i|\tilde{s}_0)p(\theta_{-i}|\tilde{s}_0) \left(p(\theta_i|\tilde{s}_1)^2 + p(\theta_{-i}|\tilde{s}_1)^2\right) \\
\Leftrightarrow &\left(p(\theta_i|\tilde{s}_1)p(\theta_i|\tilde{s}_0) - p(\theta_{-i}|\tilde{s}_1)p(\theta_{-i}|\tilde{s}_0)\right) \left(p(\theta_{-i}|\tilde{s}_1)p(\theta_i|\tilde{s}_0) - p(\theta_i|\tilde{s}_1)p(\theta_{-i}|\tilde{s}_0)\right) > 0
\end{aligned}$$

where the second equivalence is obtained using: $p(X|\tilde{s}) = p(X|\theta_i)p(\theta_i|\tilde{s}) + p(X|\theta_{-i})p(\theta_{-i}|\tilde{s})$.
Then it suffices to notice that:

- $p(\theta_i|\tilde{s}_1)p(\theta_i|\tilde{s}_0) - p(\theta_{-i}|\tilde{s}_1)p(\theta_{-i}|\tilde{s}_0) > 0 \Leftrightarrow p(\theta_i|\tilde{s}_1) > p(\theta_{-i}|\tilde{s}_0) \Leftrightarrow \mu_i > \frac{1}{2}$
- $p(\theta_{-i}|\tilde{s}_1)p(\theta_i|\tilde{s}_0) - p(\theta_i|\tilde{s}_1)p(\theta_{-i}|\tilde{s}_0) > 0 \Leftrightarrow p(\theta_i|\tilde{s}_0) > p(\theta_i|\tilde{s}_1) \Leftrightarrow i = 0$

□

Proposition 12 (Barrier to entry)**.** *If τ is a strictly increasing function, the expert has less incentives to reveal s_1 than s_0.*

Proof. We need to show that for an equal value of the private posterior $p^e(\theta_0|s_0) = p^e(\theta_1|s_1)$, condition A.4.3 is more easily met for $i = 0$ than for $i = 1$.

$p^e(\theta_0|s_0) = p^e(\theta_1|s_1)$ is equivalent to saying that the priors μ_0 and μ_1 have an equal value. We have seen that if this value is above $1/2$, the incentive compatibility condition is met in both cases and the proof is over.

Suppose now that this value is under 1/2. We first suppose the condition is met for $s = 1$, meaning:

$$\begin{aligned} p(\theta_1|s_1) &\geq a^0(\tau_0)p(\theta_0|s_1) + (1 - a^0(\tau_0))p(\theta_0|s_0) \\ &\geq a^0_{\tau_0}\left[p(\theta_0|s_1) - p(\theta_0|s_0)\right] + p(\theta_0|s_0) \\ &\geq a^1_{\tau_0}\left[p(\theta_0|s_1) - p(\theta_0|s_0)\right] + p(\theta_0|s_0) \\ &\geq a^1_{\tau_1}\left[p(\theta_0|s_1) - p(\theta_0|s_0)\right] + p(\theta_0|s_0) \end{aligned} \tag{A.4.4}$$

where the second inequality stems from $p(\theta_0|s_1) < p(\theta_0|s_0)$ and Lemma 6 which states that $a^1(\tau) > a^0(\tau) \Leftrightarrow \mu_1 < \frac{1}{2}$. The last inequality comes from the monotonicity of $a^{\tilde{s}}(.)$.

We now turn to the symmetric case $p^e(\theta_0|s_0) = p^e(\theta_1|s_1)$ (which also implies $p(\theta_0|s_1) = p^e(\theta_1|s_0)$) and inequality A.4.4 implies that the incentive compatibility condition is met for $s = 0$. □

Let us denote $\underline{\mu_1}$ and $\overline{\mu}_1$ respectively the infimum and the supremum of $IC_{\tilde{t},\tau}$. It is easy to show that $\underline{\mu_1} \leq 1/2 \leq \overline{\mu}_1$. It is also possible to show that $IC_{\tilde{t},\tau}$ is a compact set and can therefore be rewritten $\left[\underline{\mu}_1, \overline{\mu}_1\right]$.

Proposition 13 (Guru effect). *Take $\tau(.)$ increasing such that $\tau(0) = 1/2$ and $\tau(1) = 1$. Then,* $\lim_{\tilde{t}\to 1} \underline{\mu}_1 = \frac{1}{2}$.

Proof. For the first part of the proof, let us consider equation A.4.3 when $i = 1$. The left hand side is equal to $p^e(\theta_1|s_1)$ and the right hand side converges to $p^a(\theta_0|s_0)$ when $\tilde{t}$ goes to 1. But for all $\mu_1 < 1/2$, $p^a(\theta_0|s_0) > p^e(\theta_1|s_1)$ and the incentive compatible condition on this range is not satisfied when $\tilde{t}$ goes to 1. As $\mu_1 = 1/2$ always belongs to the incentive compatibility set, we deduce $\lim_{\tilde{t}\to 1} \underline{\mu}_1 = \frac{1}{2}$. □

B
Appendix — Chapter 2

Appendix B.1 provides the proofs of the propositions in Section 2.3. Appendix B.2 details the estimation strategy in the presence of conflicts of interest. Appendix B.3 gives a deeper description of our dataset. Finally, Appendix B.4 completes the exposition of our results with the exhaustive list of estimates for the reviewer-specific parameters.

B.1 Proofs of the Propositions

B.1.1 Proof of Proposition 6

We need to characterize the experts' incentive to be truthtelling when the market thinks she is. Therefore, we assume in this proof that the belief of the market on the strategy of the expert is $\tilde{\sigma} = \sigma^T$. Given this, an expert sending a recommendation r_i is perceived by the consumers as having private information $\tilde{s}_i$. The truthtelling incentive is satisfied as long as revealing s_i yields a higher expected utility:

$$\begin{aligned} &\mathbb{E}_f(u(\tilde{t})|s_i, \tilde{s}_i, t) \geq \mathbb{E}_f(u(\tilde{t})|s_i, \tilde{s}_{-i}, t) \\ \Leftrightarrow &\int_{\underline{t}}^{\bar{t}} u(\tilde{t}) dF(\tilde{t}|s_i, \tilde{s}_i, t) \geq \int_{\underline{t}}^{\bar{t}} u(\tilde{t}) dF(\tilde{t}|s_i, \tilde{s}_{-i}, t) \end{aligned} \tag{B.1.1}$$

in which $F(\tilde{t}|s_i, \tilde{s}, t)$ is the expert's expectation of her reputation update by consumers after sending a recommendation $r = \tilde{s}$.

Denoting f the density associated to F, we get:

$$f(\tilde{t}|s_i, \tilde{s}, t) = f(\tilde{t}|\tilde{s}, \theta_i) p^e(\tilde{\theta}_i|s_i, \tilde{s}, t, \mathbb{E}(\tilde{t})) + f(\tilde{t}|\tilde{s}, \theta_{-i}) p^e(\tilde{\theta}_{-i}|s_i, \tilde{s}, t, \mathbb{E}(\tilde{t}))$$

In this equation, $p^e(\tilde{\theta}|s_i, \tilde{s}, t, \mathbb{E}(\tilde{t}))$ is the expert's expectation on the consumers' belief on the state of the world at the end of the game. Notice that this expectation does not depend on the whole prior belief on the expert's ability $f(\tilde{t})$ but only on its first moment $\mathbb{E}(\tilde{t})$. Indeed, $f(\tilde{t})$ only affect $p^e(\tilde{\theta}|s_i, \tilde{s}, t)$ through the consumers' intermediate posterior after observing the recommendation of the expert: $p(\theta_i|\tilde{s}_i) = \frac{\mathbb{E}(\tilde{t})p(\theta_i)}{\mathbb{E}(\tilde{t})p(\theta_i)+(1-\mathbb{E}(\tilde{t}))p(\theta_{-i})}$.

The reputation conditional on $\tilde{s}$ and θ_i is updated as follows:

$$\begin{cases} f(\tilde{t}|\tilde{s}_i, \theta_i) = \frac{p(\tilde{s}_i|\theta_i,\tilde{t})f(\tilde{t}|\theta_i)}{\int_{\underline{t}}^{\bar{t}} p(\tilde{s}_i|\theta_i,\tilde{t})f(\tilde{t}|\theta_i)d\tilde{t}} = \frac{\tilde{t}f(\tilde{t})}{\mathbb{E}(\tilde{t})} \\ f(\tilde{t}|\tilde{s}_{-i}, \theta_i) = \frac{(1-\tilde{t})f(\tilde{t})}{1-\mathbb{E}(\tilde{t})} \end{cases}$$

Substituting these expressions in inequality B.1.1 yields:

$$\begin{aligned} \int_{\underline{t}}^{\bar{t}} u(\tilde{t}) \left(\frac{\tilde{t}f(\tilde{t})}{\mathbb{E}(\tilde{t})} - \frac{(1-\tilde{t})f(\tilde{t})}{1-\mathbb{E}(\tilde{t})} \right) \left[p^e(\tilde{\theta}_i|s_i, \tilde{s}_i, t, \mathbb{E}(\tilde{t})) - p^e(\tilde{\theta}_{-i}|s_i, \tilde{s}_{-i}, t, \mathbb{E}(\tilde{t})) \right] d\tilde{t} \geq 0 \\ \Leftrightarrow \frac{p^e(\tilde{\theta}_i|s_i, \tilde{s}_i, t, \mathbb{E}(\tilde{t})) - p^e(\tilde{\theta}_{-i}|s_i, \tilde{s}_{-i}, t, \mathbb{E}(\tilde{t}))}{\mathbb{E}(\tilde{t})(1-\mathbb{E}(\tilde{t}))} \int_{\underline{t}}^{\bar{t}} \tilde{t}\,(\tilde{t}-\mathbb{E}(\tilde{t}))\,f(\tilde{t})d\tilde{t} \geq 0 \\ \Leftrightarrow p^e(\tilde{\theta}_i|s_i, \tilde{s}_i, t, \mathbb{E}(\tilde{t})) - p^e(\tilde{\theta}_{-i}|s_i, \tilde{s}_{-i}, t, \mathbb{E}(\tilde{t})) \geq 0 \end{aligned}$$

in which the second inequality follows from $u(\tilde{t}) = \tilde{t}$ and the last one is given by:

$$\int_{\underline{t}}^{\bar{t}} \tilde{t}\,(\tilde{t}-\mathbb{E}(\tilde{t}))\,f(\tilde{t})d\tilde{t} = \int_{\underline{t}}^{\bar{t}} \tilde{t}^2 f(\tilde{t})d\tilde{t} - \mathbb{E}(\tilde{t}) \int_{\underline{t}}^{\bar{t}} \tilde{t}f(\tilde{t})d\tilde{t} = \mathbb{E}(\tilde{t}^2) - \mathbb{E}(\tilde{t})^2 \geq 0$$

B.1.2 Proof of Corollary 6

When $\tau(v) = 1\, \forall v$, $p^e(\tilde{\theta}|s, \tilde{s}, t, \mathbb{E}(\tilde{t})) = p(\theta|s, t)$, therefore the characterization is equivalent to $p(\theta_i|s_i, t) \geq p(\theta_{-i}|s_i, t) \Leftrightarrow \mu \in [1-t, t]$.

B.1.3 Proof of Proposition 7

Let us first rewrite the characterization of the incentive compatibility condition. To simplify notations, we omit t and $\mathbb{E}(\tilde{t})$ in $p^e(\tilde{\theta}_i|s_i, \tilde{s}_i, t, \mathbb{E}(\tilde{t}))$:

$$p^e(\tilde{\theta}_i|s_i, \tilde{s}_i) \geq p^e(\tilde{\theta}_{-i}|s_i, \tilde{s}_{-i})$$

$$\Leftrightarrow p(X_i|s_i, \tilde{s}_i)p(\tilde{\theta}_i|X_i, \tilde{s}_i) + p(X_{-i}|s_i, \tilde{s}_i)p(\tilde{\theta}_i|X_{-i}, \tilde{s}_i) \geq \\ p(X_i|s_i, \tilde{s}_{-i})p(\tilde{\theta}_{-i}|X_i, \tilde{s}_{-i}) + p(X_{-i}|s_i, \tilde{s}_{-i})p(\tilde{\theta}_{-i}|X_{-i}, \tilde{s}_{-i})$$

$$\Leftrightarrow \left[(2\tau(\nu(\tilde{s}_i)) - 1)p(\theta_i|s_i, t) + 1 - \tau(\nu(\tilde{s}_i))\right] p(\tilde{\theta}_i|X_i, \tilde{s}_i) + \\ \left[(1 - 2\tau(\nu(\tilde{s}_i)))p(\theta_i|s_i, t) + \tau(\nu(\tilde{s}_i))\right] p(\tilde{\theta}_i|X_{-i}, \tilde{s}_i) \geq \\ \left[(2\tau(\nu(\tilde{s}_{-i})) - 1)p(\theta_i|s_i, t) + 1 - \tau(\nu(\tilde{s}_{-i}))\right] p(\tilde{\theta}_{-i}|X_i, \tilde{s}_{-i}) + \\ \left[(1 - 2\tau(\nu(\tilde{s}_{-i})))p(\theta_i|s_i, t) + \tau(\nu(\tilde{s}_{-i}))\right] p(\tilde{\theta}_{-i}|X_{-i}, \tilde{s}_{-i})$$

$$\Leftrightarrow p(\theta_i|s_i, t)\left\{\left[2\tau(\nu(\tilde{s}_i)) - 1\right]\left(p(\tilde{\theta}_i|X_i, \tilde{s}_i) - p(\tilde{\theta}_i|X_{-i}, \tilde{s}_i)\right) + \\ \left[2\tau(\nu(\tilde{s}_{-i})) - 1\right]\left(p(\tilde{\theta}_{-i}|X_{-i}, \tilde{s}_{-i}) - p(\tilde{\theta}_{-i}|X_i, \tilde{s}_{-i})\right)\right\} \geq \\ (1 - \tau(\nu(\tilde{s}_{-i})))p(\tilde{\theta}_{-i}|X_i, \tilde{s}_{-i}) + \tau(\nu(\tilde{s}_{-i}))p(\tilde{\theta}_{-i}|X_{-i}, \tilde{s}_{-i}) - \\ (1 - \tau(\nu(\tilde{s}_i)))p(\tilde{\theta}_i|X_i, \tilde{s}_i) - \tau(\nu(\tilde{s}_i))p(\tilde{\theta}_i|X_{-i}, \tilde{s}_i)$$

Noticing that the right-hand side does not depend on the true ability of the reviewer, and that the expression between brackets in the left-hand side is positive and does not depend on the true ability, we get that the inequality is more slack when $p(\theta_i|s_i, t)$ increases. Therefore, if the incentive compatibility is satisfied for a true ability $t^a \leq t^b$, it is also satisfied for t^b.

B.2 Identification and Estimation Strategy of the Model with Potential Conflicts of Interest

As highlighted in DellaVigna and Hermle [15], potential conflicts of interest could arise in movie reviews if the reviewer's newspaper and the movie's production company belong to the same media outlet. An example would be a movie released by the 20^{th} Century Fox reviewed in the New York Post, both members of the media outlet Newscorp. We detail all possible conflicts of interest in the data section.

We include these conflict of interest in our structural estimation by modelling the behavior of the movie reviewer facing a movie from her own media outlet. We denote b^c her bias caused by the conflict of interest and say that $b^c = p(r_1|\hat{r_0})$ in which $\hat{r_0}$ is the recommendation the reviewer would have sent in the absence of any conflict of interest. This is therefore equivalent to an alternative model in which the movie reviewer plays the game as described in the main theoretical model but switches her recommendation with probability b^c if she were to send a negative recommendation concerning a movie from her media outlet.

The identification and estimation strategies are similar to those for the main model except for the reduced-form conditional grading probabilities, which change in the following way in case of a conflict of interest:

Table B.2.1: Probabilities of giving a good review conditional on the true state of the world and on the prior in the case of a conflict of interest.

Prior	$\gamma^c_{i,1}$	$\gamma^c_{i,0}$
$\mu_j < \underline{\mu_i}$	$(1-t)b^c + t(1-b^- + b^- b^c)$	$tb^c + (1-t)(1-b^- + b^- b^c)$
$\mu_j \in \left[\underline{\mu_i}, \overline{\mu_i}\right]$	$t + (1-t)b^c$	$(1-t) + tb^c$
$\mu_j > \overline{\mu_i}$	$t + (1-t)(b^+ + (1-b^+)b^c)$	$(1-t) + t(b^+ + (1-b^+)b^c)$

We estimate this model by maximizing the following maximum likelihood:

$$\max_{\{t_i, b_i^-, b_i^+,\}_{i=1}^n, \beta, \delta, b^c} \sum_j \log \Bigg[\mu(\omega_j;\beta) \prod_{i=1}^n \Big\{ \left(\gamma^c_{i,1}{}^{r_i}(1-\gamma^c_{i,1})^{1-r_i}\right)^{C_{ij}} \left(\gamma_{i,1}{}^{r_i}(1-\gamma_{i,1})^{1-r_i}\right)^{1-C_{ij}} \Big\}$$
$$+(1-\mu(\omega_j;\beta)) \prod_{i=1}^n \Big\{ \left(\gamma^c_{i,1}{}^{r_i}(1-\gamma^c_{i,1})^{1-r_i}\right)^{C_{ij}} \left(\gamma_{i,1}{}^{r_i}(1-\gamma_{i,1})^{1-r_i}\right)^{1-C_{ij}} \Big\} \Bigg]$$

$$\text{s.t. } t_i \in \left[\frac{1}{2}, 1\right], \quad b_i^+, b_i^-, b^c \in [0,1], \delta \geq 0 \qquad \text{(B.2.1)}$$

In this likelihood, C_{ij} is a dummy variable taking value 1 if there is a potential conflict of interest between movie reviewer *i* for movie *j*. $\gamma^c_{i,\theta}$ and $\gamma_{i,\theta}$ in this expression are taken from Tables B.2.1 and 2.4.1, their value depending on the prior. For instance, if $\mu(\omega_j;\beta) < \underline{\varphi}(\mathbb{E}(t_i|\omega_i;\delta))$, the expressions for $\gamma^c_{i,\theta}$ and $\gamma_{i,\theta}$ are taken in the first line of Tables B.2.1 and 2.4.1. Notice also that we do not estimate an impact of conflicts of interest which is specific to reviewers because of a lack of observations.

B.3 Description of the Data

Table B.3.1: Variables used in our estimations, their source, and description.

Variable	**Description**	**Source**
Title	Title of the movie	
Director	Name of the movie's director	
Year	Year of release of the movie in the USA	*www.imdb.com*

Variable	Description	Source
r_{ij}	The recommendation or review given by movie reviewer *i* to movie *j*, either 1 (fresh) or 0 (rotten)	*www.rottentomatoes.com*
Reviewer	Name of the author of the review	*www.rottentomatoes.com*
Newspaper	Name of the medium in which the review is published, which is not necessarily a newspaper	*www.rottentomatoes.com*
Date of the review		*www.rottentomatoes.com*
Number of previously directed films	Number of previous films by the director. In the case of several directors for one movie, the sum of the number of previous films.	*www.imdb.com*
Budget	The estimated production budget of a movie in 2013's US dollars. We used historical data on exchange rates for non-US movies and the Consumer Price Index (CPI-U) data provided by the U.S. Department of Labor Bureau of Labor Statistic to take into account inflation.	*www.imdb.com*
Box-Office US	Gross Profit in the USA, expressed in 2013's US dollars	*www.imdb.com*
USA	Dummy equal to 1 if the movie is produced in the USA only.	*www.imdb.com*
Coproduction USA	Dummy equal to 1 if the movie is produced in the USA and at least another country.	*www.imdb.com*
Genre	A proxy for the type of the movie (e.g. action, thriller, documentary)	*www.imdb.com*
G, PG, PG-13, R, NC-17	MPAA rating of the movie	*www.imdb.com*
Remake	Dummy equal to 1 if the movie is a remake.	*www.imdb.com*
Sequel	Dummy equal to 1 if the movie is a sequel.	*www.imdb.com*
Google Search Index	Index taking a value 100 the month Peter Travers, our benchmark reviewer, got the most searches of his name on Google. All other values are computed related to this reference point. Note that these values are updated each day. We collected our data on reputation on May 6th, 2013.	*www.google.com/trends*
Production Company	Main production company of the movie, in general the distributor	*www.metacritic.com*, *www.rottentomatoes.com*

Variable	Description	Source
Conflict of Interest	Dummy equal to 1 if Newspaper and Production Company belong to the same media outlet	DellaVigna and Hermle [15], *www.wikipedia.com*
Format	Format of the medium through which the review is published (e.g. newspaper, magazine, tabloid)	*www.wikipedia.com*
Content	Content of the medium through which the review is published (e.g. news, culture, cinema)	*www.wikipedia.com*
Target	Commercial target of the medium through which the review is published, i.e. whether it is designed for a general audience or. professionals.	*www.wikipedia.com*

Table B.3.2: Movie reviewers included in the empirical analysis along with their number of reviews in our sample, their average Google Search index through the sampling period, and the media through which they published.

Reviewer's Name	Total # Reviews	Avg. Google Search index	Media
Ann Hornaday	857	0.88	Washington Post
Ao Scott	1,235	6	At the Movies, New York Times
Carrie Rickey	1,020	0.045	Philadelphia Inquirer
Claudia Puig	1,469	0.67	USA Today
Colin Covert	1,435	0.1	Chicago Tribune, Minneapolis Star Tribune
David Edelstein	1,029	11.15	NPR, New York Magazine, Slate
Desson Thomson	1,266	0	Washington Post
Elizabeth Weitzman	997	0.09	New York Daily News
James Berardinelli	2,858	26.36	Reelviews
Joe Baltake	904	0	Passionate Moviegoer, Sacramento Bee
Kenneth Turan	1,124	3.89	Los Angeles Times, Newsday
Kirk Honeycutt	1,163	0.009	Hollywood Reporter
Kyle Smith	1,010	40.41	New york Post
Liam Lacey	998	0.2	Globe and Mail
Lisa Schwarzbaum	1,633	0.5	Entertainment Weekly
Lou Lumenick	1,723	0.63	New York Post

Reviewer's Name	**Total # Reviews**	**Avg. Google Search index**	**Media**
Michael Phillips	1,157	72.62	At the Movies, Chicago Tribune
Mick Lasalle	1,579	2.62	Houston Chronicle, San Francisco Chronicle
Moira Macdonald	1,415	0.18	Seattle Times
Owen Gleiberman	1,971	0.79	CNN.com, Entertainment Weekly
Peter Howell	1,260	9.75	Toronto Star
Peter Travers	1,879	23.24	Rolling Stone
Rex Reed	940	18.125	New York Observer
Richard Roeper	1,736	35.31	Chicago Sun-Times, Ebert & Roeper, Richard Roeper.com
Robert Denerstein	946	0	Denver Rocky Mountain
Roger Ebert	3,277	1100	Chicago Sun-Times, Denver Post, Detroit News, Ebert & Roeper
Stephanie Zacharek	944	1.21	CNN.com, Film.com, Los Angeles Times, NPR, Salon.com, Village Voice
Stephen Holden	1,050	5.17	New York Times
Stephen Whitty	1,403	0	Newark Star-Ledger
Steven Rea	1,232	1.18	Philadelphia Inquirer
Terry Lawson	1,193	3.34	Detroit Free Press, Miami Herald
Todd Mccarthy	903	3.125	Hollywood Reporter, Variety, indieWire
Ty Burr	1,313	1.375	Boston Globe, Dallas Morning News, Entertainment Weekly
Unknown Reviewer	2,142	0	Time Out (40%)
Wesley Morris	1,217	3.49	Boston Globe

Table B.3.3: Potential Conflicts of Interest

Media Outlet	Production Companies	Media	Years of Interactions	Concerned Reviewers	# Observations
News Corp.	20th Century Fox, Fox Searchlight Pictures	New York Post	1993-2013	Kyle Smith	56
				Lou Lumenick	74
Time Warner	Warner Bros Pictures, Picturehouse, HBO	Entertainment Weekly	1990-2013	Lisa Schwarzbaum	88
				Owen Gleiberman	102
				Unknown Reviewer	3
		TIME Magazine	1990-2013	Unknown Reviewer	1
	New Line	Entertainment Weekly	1996-2010	Lisa Schwarzbaum	18
				Owen Gleiberman	31
		TIME Magazine	1996-2010	Unknown Reviewer	1
	Fine Line Features	Entertainment Weekly	1996-2013	Lisa Schwarzbaum	4
				Owen Gleiberman	3
Disney Media Group	Walt Disney Pictures, Buena Vista	At the Movies	2007-2010	Ao Scott	5
				Michael Phillips	1
		Ebert & Roeper	2007-2013	Richard Roeper	40
	Miramax	At the Movies	2007-2010	Ao Scott	2
				Michael Phillips	1
		Ebert & Roeper	2007-2010	Richard Roeper	22

B.4 Additional Tables

Table B.4.1: ML estimates for reviewer-specific parameters — Specification (II) includes the budget and conflicts of interest

	(I)						(II)					
	t		b^-		b^+		t		b^-		b^+	
Reviewer's Name	Coeff.	SE	Coeff.	SE	Coeff.	SE	Coeff.	SE	Coeff.	SE	Coeff.	SE
Ann Hornaday	0.781	(0.021)	0.027	(0.045)	0.094	(0.047)	0.785	(0.029)	0.082	(0.059)	0.143	(0.063)
Ao Scott	0.758	(0.017)	0.001	(0.034)	0.175	(0.051)	0.741	(0.024)	0.032	(0.051)	0.089	(0.056)
Carrie Rickey	0.784	(0.015)	0.000	(0.000)	0.367	(0.066)	0.767	(0.027)	0.000	(0.015)	0.349	(0.062)
Claudia Puig	0.822	(0.006)	0.073	(0.038)	0.195	(0.041)	0.833	(0.018)	0.124	(0.047)	0.140	(0.041)
Colin Covert	0.765	(0.009)	0.002	(0.000)	0.184	(0.043)	0.773	(0.018)	0.002	(0.001)	0.235	(0.050)
David Edelstein	0.769	(0.014)	0.001	(0.030)	0.250	(0.055)	0.773	(0.025)	0.062	(0.044)	0.223	(0.057)
Desson Thomson	0.821	(0.023)	0.188	(0.072)	0.001	(0.028)	0.819	(0.029)	0.155	(0.083)	0.001	(0.031)
Elizabeth Weitzman	0.804	(0.015)	0.024	(0.051)	0.065	(0.041)	0.815	(0.024)	0.081	(0.061)	0.075	(0.047)
James Berardinelli	0.768	(0.011)	0.155	(0.040)	0.110	(0.030)	0.772	(0.019)	0.172	(0.043)	0.077	(0.046)
Joe Baltake	0.657	(0.043)	0.000	(0.000)	0.432	(0.090)	0.663	(0.052)	0.000	(0.004)	0.510	(0.101)
Kenneth Turan	0.789	(0.017)	0.007	(0.035)	0.233	(0.060)	0.758	(0.033)	0.030	(0.046)	0.217	(0.086)
Kirk Honeycutt	0.734	(0.014)	0.029	(0.044)	0.095	(0.051)	0.734	(0.026)	0.039	(0.044)	0.045	(0.048)
Kyle Smith	0.619	(0.015)	0.345	(0.047)	0.000	(0.000)	0.654	(0.026)	0.360	(0.065)	0.000	(0.000)
Liam Lacey	0.790	(0.026)	0.406	(0.055)	0.004	(0.000)	0.794	(0.029)	0.439	(0.065)	0.004	(0.006)
Lisa Schwarzbaum	0.823	(0.009)	0.000	(0.030)	0.164	(0.056)	0.823	(0.024)	0.016	(0.042)	0.132	(0.048)
Lou Lumenick	0.821	(0.006)	0.202	(0.044)	0.046	(0.025)	0.823	(0.021)	0.178	(0.060)	0.003	(0.042)
Michael Phillips	0.799	(0.006)	0.092	(0.035)	0.091	(0.037)	0.789	(0.023)	0.111	(0.053)	0.075	(0.050)
Mick Lasalle	0.665	(0.011)	0.018	(0.039)	0.147	(0.050)	0.661	(0.024)	0.003	(0.035)	0.119	(0.051)
Moira Macdonald	0.838	(0.000)	0.003	(0.009)	0.241	(0.052)	0.836	(0.021)	0.003	(0.036)	0.213	(0.053)
Owen Gleiberman	0.768	(0.011)	0.000	(0.000)	0.263	(0.038)	0.770	(0.025)	0.000	(0.008)	0.267	(0.049)
Peter Howell	0.807	(0.014)	0.008	(0.035)	0.234	(0.049)	0.811	(0.022)	0.045	(0.043)	0.273	(0.054)
Peter Travers	0.877	(0.006)	0.052	(0.028)	0.357	(0.046)	0.857	(0.019)	0.059	(0.038)	0.265	(0.052)
Rex Reed	0.670	(0.016)	0.129	(0.068)	0.001	(0.027)	0.680	(0.029)	0.224	(0.075)	0.001	(0.018)
Richard Roeper	0.791	(0.009)	0.008	(0.006)	0.355	(0.044)	0.787	(0.020)	0.004	(0.015)	0.266	(0.053)
Robert Denerstein	0.898	(0.009)	0.107	(0.035)	0.001	(0.034)	0.898	(0.021)	0.146	(0.062)	0.013	(0.042)
Roger Ebert	0.782	(0.000)	0.000	(0.000)	0.365	(0.035)	0.775	(0.015)	0.000	(0.007)	0.353	(0.047)
Stephanie Zacharek	0.700	(0.020)	0.080	(0.048)	0.004	(0.011)	0.688	(0.028)	0.115	(0.062)	0.004	(0.013)
Stephen Holden	0.790	(0.020)	0.057	(0.051)	0.102	(0.040)	0.800	(0.032)	0.076	(0.085)	0.014	(0.052)
Stephen Whitty	0.722	(0.006)	0.255	(0.053)	0.004	(0.000)	0.726	(0.022)	0.276	(0.052)	0.004	(0.002)
Steven Rea	0.771	(0.011)	0.023	(0.041)	0.227	(0.056)	0.763	(0.025)	0.164	(0.065)	0.244	(0.055)
Terry Lawson	0.799	(0.020)	0.104	(0.050)	0.186	(0.053)	0.795	(0.029)	0.151	(0.067)	0.147	(0.060)
Todd Mccarthy	0.782	(0.012)	0.039	(0.039)	0.109	(0.047)	0.790	(0.028)	0.019	(0.054)	0.095	(0.061)
Ty Burr	0.791	(0.009)	0.047	(0.034)	0.000	(0.021)	0.772	(0.022)	0.051	(0.045)	0.000	(0.026)
Unknown Reviewer	0.749	(0.016)	0.080	(0.056)	0.055	(0.051)	0.757	(0.039)	0.217	(0.102)	0.109	(0.073)
Wesley Morris	0.738	(0.006)	0.318	(0.055)	0.005	(0.000)	0.742	(0.024)	0.330	(0.054)	0.001	(0.000)
# Observations:	30440						22674					

Notes: Bootstrap Standard Errors are computed on 100 iterations for (I) and 500 for (II)

Table B.4.2: ML Estimates for Movie Specific Parameters on a Subset of the Data Set — Specification (II) is defined as in section 2.6; Specification (III) excludes strategic biases; Specification (IV) is a naive logit estimation of the prior.

	(II)		(III)		(IV)	
	Coeff.	SE	Coeff.	SE	Coeff.	SE
Movie Specific, β:						
Constant	1.124	(0.010)	3.076	(0.842)	2.208	(0.213)
Origin: USA	-1.440	(0.037)	-0.834	(0.226)	-0.433	(0.056)
Origin: co-production USA	-1.556	(0.102)	-0.823	(0.244)	-0.428	(0.059)
Remake	-0.440	(0.232)	-0.594	(0.270)	-0.256	(0.062)
Sequel	-0.434	(0.140)	-0.263	(0.210)	-0.076	(0.049)
Director's number of previous movies	0.016	(0.000)	0.029	(0.010)	0.017	(0.002)
G rating	0.953	(3.439)	1.704	(1.224)	0.972	(0.136)
PG rating	0.470	(0.201)	0.151	(0.206)	0.151	(0.049)
R rating	0.530	(0.060)	0.361	(0.129)	0.219	(0.035)
NC-17 rating	0.345	(6.490)	0.077	(1.976)	0.165	(0.222)
log(budget)	0.000	(0.000)	-0.157	(0.050)	-0.111	(0.013)
Reputation, δ:						
Google Search Index	3.5×10^{-4}	(0.000)				
Conflict of interest, b^c:						
Average Bias	0.043	(0.040)				
Likelihood	-10422		-10620		-12144	
# Observations:	17882		17882		17882	

Notes: Standard Errors are bootstrapped on 189 iterations for (II) and 200 for (III).

Table B.4.3: ML Estimates for Reviewer-Specific Parameters on a Subset of the Data Set — Specification (II) is defined as in section 2.6; Specification (III) excludes strategic biases.

	(II)						(III)	
	t		b^-		b^+		t	
Reviewer's Name	Coeff.	SE	Coeff.	SE	Coeff.	SE	Coeff.	SE
Ann Hornaday	0.782	(0.029)	0.125	(0.067)	0.205	(0.066)	0.768	(0.023)
Ao Scott	0.756	(0.028)	0.020	(0.048)	0.097	(0.060)	0.749	(0.022)
Carrie Rickey	0.760	(0.029)	0.000	(0.014)	0.296	(0.083)	0.754	(0.023)
Claudia Puig	0.822	(0.020)	0.108	(0.051)	0.128	(0.051)	0.805	(0.014)
Colin Covert	0.768	(0.020)	0.000	(0.001)	0.224	(0.051)	0.746	(0.017)
David Edelstein	0.775	(0.025)	0.039	(0.046)	0.207	(0.060)	0.770	(0.021)
Desson Thomson	0.806	(0.032)	0.154	(0.093)	0.000	(0.021)	0.804	(0.027)
Elizabeth Weitzman	0.792	(0.030)	0.084	(0.074)	0.093	(0.064)	0.778	(0.021)
James Berardinelli	0.784	(0.017)	0.164	(0.052)	0.098	(0.048)	0.762	(0.015)
Joe Baltake	0.644	(0.062)	0.000	(0.011)	0.515	(0.106)	0.652	(0.045)
Kenneth Turan	0.769	(0.032)	0.027	(0.056)	0.215	(0.090)	0.768	(0.026)
Kirk Honeycutt	0.727	(0.026)	0.000	(0.048)	0.071	(0.061)	0.717	(0.023)
Kyle Smith	0.646	(0.028)	0.361	(0.071)	0.000	(0.000)	0.630	(0.025)
Liam Lacey	0.788	(0.029)	0.461	(0.076)	0.000	(0.011)	0.744	(0.025)
Lisa Schwarzbaum	0.822	(0.022)	0.050	(0.053)	0.111	(0.060)	0.802	(0.019)
Lou Lumenick	0.825	(0.023)	0.170	(0.071)	0.000	(0.047)	0.799	(0.018)
Michael Phillips	0.796	(0.021)	0.122	(0.056)	0.070	(0.053)	0.783	(0.016)
Mick Lasalle	0.655	(0.022)	0.000	(0.038)	0.123	(0.068)	0.657	(0.022)
Moira Macdonald	0.821	(0.024)	0.000	(0.025)	0.207	(0.063)	0.807	(0.018)
Owen Gleiberman	0.766	(0.025)	0.000	(0.024)	0.290	(0.054)	0.738	(0.024)
Peter Howell	0.807	(0.025)	0.042	(0.047)	0.263	(0.060)	0.769	(0.020)
Peter Travers	0.849	(0.022)	0.078	(0.051)	0.262	(0.059)	0.825	(0.016)
Rex Reed	0.670	(0.030)	0.280	(0.081)	0.000	(0.001)	0.620	(0.028)
Richard Roeper	0.785	(0.022)	0.000	(0.008)	0.257	(0.059)	0.749	(0.019)
Robert Denerstein	0.898	(0.022)	0.201	(0.069)	0.019	(0.043)	0.848	(0.024)
Roger Ebert	0.778	(0.020)	0.000	(0.010)	0.337	(0.054)	0.761	(0.015)
Stephanie Zacharek	0.683	(0.027)	0.078	(0.071)	0.000	(0.002)	0.656	(0.027)
Stephen Holden	0.783	(0.036)	0.088	(0.090)	0.026	(0.056)	0.793	(0.028)
Stephen Whitty	0.698	(0.020)	0.297	(0.070)	0.000	(0.000)	0.670	(0.019)
Steven Rea	0.765	(0.029)	0.118	(0.070)	0.226	(0.055)	0.740	(0.022)
Terry Lawson	0.761	(0.032)	0.146	(0.080)	0.200	(0.069)	0.737	(0.023)
Todd Mccarthy	0.791	(0.030)	0.000	(0.051)	0.106	(0.071)	0.780	(0.025)
Ty Burr	0.782	(0.021)	0.051	(0.061)	0.000	(0.032)	0.769	(0.019)
Unknown critic	0.759	(0.041)	0.237	(0.108)	0.101	(0.083)	0.750	(0.033)
Wesley Morris	0.744	(0.028)	0.335	(0.065)	0.000	(0.000)	0.717	(0.021)
# Observations:	17882						17882	

Notes: Bootstrap Standard Errors are computed on 189 iterations for (II) and 200 for (III).

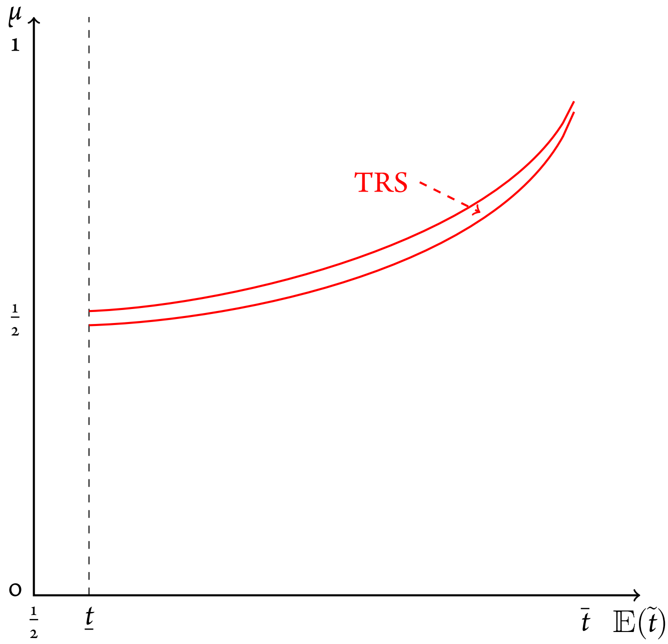

Figure B.4.1: Truthful revelation set for $\underline{t} = 0.55$ and $\tau(v) = \frac{1+v}{2}$.

C
Appendix — Chapter 3

C.1 Long-run Consumers

Proof of Lemma 1: From Table 3.4.1, which summarizes the discounted payoffs of the incumbency game, it is easy to see that whatever the equilibrium prices, the consumers obtain higher payoffs when they coordinate and take the same action, that is, either both consumers stay or switch (St, St) and (Sw, Sw). However, absent coordination, the decision of a consumer depends on her belief about the action of the other consumer. We start from the case where a consumer is considering switching, and, we denote by μ her belief on the other consumer also switching. Therefore, given this belief, a consumer decides to switch if:

$$\frac{1-p^I}{1-\delta} \leq \mu \times \left(1 - p^E - S + \delta\frac{1-p^I}{1-\delta}\right) + (1-\mu) \times \left(1 - p^E - S\right).$$

By solving this expression for the price of the entrant, we obtain:

$$p^E(\mu) = \frac{(1-\mu\delta)p^I - (1-\mu)\delta - (1-\delta)S}{1-\delta},$$

which increases on the belief that the other consumer also switches. In equilibrium, the incumbent will set a price such that the entrant cannot attract consumers without incurring negative profits:

$$2\bar{p}^E + 2\delta\frac{p^I}{1-\delta} < 0.$$

Combining both expressions, the price of the incumbent is:

$$p^I(\mu) = \frac{(1-\mu)\delta + (1-\delta)S}{1+(1-\mu)\delta}.$$

If we assume perfect coordination — $\mu = 1$ — we obtain the equilibrium prices $(p^I, p^E) = ((1-\delta)S, -\delta S)$ and this coincides with the one obtained in Biglaiser et al. [8], where the authors assume perfect coordination of consumers. Notice also that the equilibrium incumbent price is decreasing with the belief μ that the other consumer switches. And the price is minimal when there is perfect coordination of consumers on switching. Therefore, the maximum incumbency price is when consumers are pessimistic. A pessimistic consumer only switches if it is a dominant strategy to do so, and this creates extra frictions that are beneficial to the incumbent.

In the other case, a consumer considers staying with the incumbent and with a probability μ he believes that the other consumer will also stay. Given this beliefs, the consumer switches to the entrant if:

$$\mu \times \frac{1-p^I}{1-\delta} + (1-\mu) \times (1-p^I) \leq 1 - p^E - S + \delta \frac{1-p^I}{1-\delta},$$

which is equivalent to:

$$p^E(\mu) = \frac{(1+\delta(\mu-2))\, p^I + (1-\mu)\delta - (1-\delta)S}{1-\delta}.$$

By applying the same reasoning, the incumbency price is given by:

$$p^I(\mu) = \frac{(1-\delta)S - (1-\mu)\delta}{1-(1-\mu)\delta},$$

which increases with μ. Here perfect coordination means that both consumers stay with the incumbent $\mu = 1$ and we obtain the same result as in Biglaiser et al. [8]. However, we establish a lower bound of the price of the incumbent when no coordination between consumers occur. Furthermore, as it is a dominant strategy to set a negative price, the incumbency price is then $p^I(\mu) = \max\left\{0, \frac{(1-\delta)S-(1-\mu)\delta}{1-(1-\mu)\delta}\right\}$. Substituting those incumbency prices by the different levels of coordination μ we obtain the multiplicity of equilibria as stated in the Lemma.

C.2 Short-run consumers

Proof of Lemma (3) : Shilony [49] proves that there exists a unique equilibrium in the game where the range of the support is $\bar{p} - \underline{p} \leq 2S$. We solve the problem by considering different ranges of p belonging to the support of the game:

$\underline{p} + S \leqslant p < \bar{p}$. With p in this range, we get that $F(min(p+S, R)) = 1$ and

$F(max(0, p-S)) = F(p-S)$. Therefore, we obtain $F(p-S) = 1 - \frac{V}{p}$ which is equivalent to:

$$F(p) = 1 - \frac{V}{p+S} \quad \text{for} \quad \underline{p} \leqslant p < \bar{p} - S.$$

$\underline{p} \leqslant p < \bar{p} - S$. With p in this range, we get that $F(max(0, p-S)) = 0$ and $F(min(p+S, 1)) = F(p+S)$. Therefore, we obtain $F(p+S) = 2 - \frac{V}{p}$ which is equivalent to:

$$F(p) = 2 - \frac{V}{p-S} \quad \text{for} \quad \underline{p} + S \leqslant p < \bar{p}.$$

$\bar{p} - S < p < \underline{p} + S$. The difference between any price in this range and any point in the support is lower than the value of switching costs. As a result, no firm loses or gains consumers when setting this price. Accordingly, no firm plays with positive probability in this range of prices. Since the distribution function is monotone and increasing with p, the distribution function in this range is:

$$F(p) = 1 - \frac{V}{\bar{p}} \quad \text{for} \quad \bar{p} - S < p < \underline{p} + S.$$

Therefore, the distribution function is:

$$F(p) = \begin{cases} 0 & \text{if } p < \underline{p}, \\ 1 - \frac{V}{p+S} & \text{if } \underline{p} \leqslant p < \bar{p} - S, \\ 1 - \frac{V}{\bar{p}} & \text{if } \bar{p} - S < p < \underline{p} + S, \\ 2 - \frac{V}{p-S} & \text{if } \underline{p} + S \leqslant p < \bar{p}, \\ 1 & \text{if } p \geqslant \bar{p}. \end{cases} \tag{C.2.1}$$

If the distribution function in expression (C.2.1) constitutes a mixed strategy equilibrium, then the expected profit of any firm is the same regardless at which point of the support it plays, that is, $\pi_i(p_i, \sigma_j^*) = V \, \forall p_i \in [\underline{p}, \bar{p}]$.

Consider that firm i sets a price $p_i = \bar{p} - S - \varepsilon$ for any $\varepsilon \in (0, \bar{p} - \underline{p} - S]$. By setting this price, the firm never loses its captive consumer and attracts the consumer of the other firm when it charges a price $p_j \geq \bar{p} - \varepsilon$. Hence, the expected profit for firm i is:

$$(\bar{p} - S - \varepsilon)\left(1 + [1 - F(\bar{p} - \varepsilon)]\right) = (\bar{p} - S - \varepsilon) \times \left(\frac{V}{\bar{p} - S - \varepsilon}\right) = V.$$

Consider that firm i sets a price $p_i = \underline{p} + S + \varepsilon$ for any $\varepsilon \in (0, \bar{p} - \underline{p} - S]$. In this case, it will lose its consumer if the other firm sets a price $p_j \leq \underline{p} + \varepsilon$. Hence the expected profit of the firm is:

$$(\underline{p} + S + \varepsilon)\left(1 - 1 \times \left[F(\underline{p} + \varepsilon)\right]\right) = \left(\underline{p} + S + \varepsilon\right) \times \left(\frac{V}{\underline{p} + S + \varepsilon}\right) = V.$$

Furthermore, if no firm plays with a positive probability the extremes of the support i.e $A(\underline{p}) = 0$ and $A(\bar{p}) = 0$, then no firm i loses its captive consumer nor gains the consumer of the rival by setting a price $p_i = \underline{p} + S$ or $p_i = \bar{p} - S$ respectively.[1]

If $p_i = \underline{p} + S$, the firm gets $\underline{p} + S$ and this should be equal to V. This gives the condition:

$$\underline{p} = V - S. \tag{C.2.2}$$

We obtain the same condition if and only if the distribution function in (C.2.1) has no atom at the lower bound of the support, $\lim_{p \downarrow \underline{p}} F(\underline{p}) = 1 - \frac{V}{\underline{p}+S} = 0$.

If $p_i = \bar{p} - S$, the firm gets $\bar{p} - S$ and this should be equal to V. This gives the condition:

$$\bar{p} = V + S. \tag{C.2.3}$$

The same condition is obtained if and only if the distribution function in (C.2.1) has no atom at the upper bound of the support, $\lim_{p \uparrow \bar{p}} F(p) = 2 - \frac{V}{\bar{p}-S} = 1$

Combining condition (C.2.2) and (C.2.3), we obtain that $\bar{p} - \underline{p} = 2S$ and the expected profit of the firm is $V = \underline{p} + S$. As a result, expression (C.2.1) can be written only as a function of the lower bound of the support $\underline{p}$.

$$F(p) = \begin{cases} 0 & \text{if } p < \underline{p}, \\ 1 - \frac{\underline{p}+S}{p+S} & \text{if } \underline{p} \leqslant p < \underline{p} + S, \\ 2 - \frac{\underline{p}+S}{p-S} & \text{if } \underline{p} + S \leqslant p < \underline{p} + 2S, \\ 1 & \text{if } p \geqslant \underline{p} + 2S. \end{cases} \tag{C.2.4}$$

To obtain the final expression of the distribution function, we make explicit use of the following Lemma:

Lemma 7. *The distribution function F(p) in expression (C.2.4) is atomless.*

Proof. We proceed by construction. First, using our tie-breaking rule, we calculate the profits that any firm obtains if it plays at the upper bound of the support. This allows us to obtain the value for the lower bound of the support. Finally, we use this value to prove that no atom exists at $\underline{p} + S$. If a firm sets a price equal to $p = \underline{p} + 2S.$, then, by the tie-breaking rule, the firm loses consumers if the other firm sets a price equal or below $\underline{p} + S$ and this happens with probability $\frac{\underline{p}-S}{\underline{p}}$. Therefore, by setting this price, the firm obtains profits:

$$\pi_i(\underline{p} + 2S, \sigma_j) = (\underline{p} + 2S) \times \left(1 - \frac{\underline{p} - S}{\underline{p}}\right) \iff (\underline{p} + 2S) \times \frac{S}{\underline{p}} = V$$
$$\iff (\underline{p} + 2S) \times \frac{S}{\underline{p}} = \underline{p} + S \iff \underline{p} = \sqrt{2}S$$

[1]Firms only lose or gain consumers by setting this price if firms played with a positive probability at the extremes of the support.

□

Finally, with the lower bound $\underline{p} = \sqrt{2}S$, there exists no atom at $\underline{p} + S$. From the picture above, we see that the atom is equal to $A(\underline{p} + S) = \frac{\underline{p}-S}{\underline{p}} - \frac{S}{\underline{p}+2S}$ and it is easy to check that $A(\underline{p} + S) = 0$ when $\underline{p} = \sqrt{2}S$ and the distribution function is equal to:

$$\sigma^{II} = F(p) = \begin{cases} 0 & \text{if } p < \sqrt{2}S, \\ 1 - \frac{(1+\sqrt{2})S}{p+S} & \text{if } \sqrt{2}S \leq p < (1+\sqrt{2})\, S, \\ 2 - \frac{(1+\sqrt{2})S}{p-S} & \text{if } (1+\sqrt{2})\, S \leq p < (2+\sqrt{2})\, S, \\ 1 & \text{if } p \geq (2+\sqrt{2})\, S, \end{cases} \tag{C.2.5}$$

and the expected profit of firms playing this equilibrium is $V = (1+\sqrt{2})\, S$ which is the one introduced in the proposition. (ii) If switching costs are $S \in \left(\frac{1}{2+\sqrt{2}}, \frac{1}{2}\right)$, there are two symmetric mixed strategy equilibria where firms randomize over a support $[\underline{p}, 1]$ according to $F(p)$ given by

(a) and each firm's expected profit is equals to:

$$V = 1 - S$$

(b) and each firm's expected profit is equals to:

$$V = \frac{S + \sqrt{S(4+S)}}{2}$$

From expression (C.2.5), the upper-bound is an increasing function of the level of switching costs S. Since no firm sets a price above the consumers reservation price, the distribution function we obtained is only valid for a value of switching costs $S \in \left(0, \frac{1}{2+\sqrt{2}}\right]$. Consequently, when $S \in \left(\frac{1}{2+\sqrt{2}}, \frac{1}{2}\right)$, the upper-bound of the support is equal to 1 and the distribution function is:

$$F(p) = \begin{cases} 0 & \text{if } p < \underline{p}, \\ 1 - \frac{V}{p+S} & \text{if } \underline{p} \leqslant p < 1 - S, \\ 1 - V & \text{if } 1 - S < p < \underline{p} + S, \\ 2 - \frac{V}{p-S} & \text{if } \underline{p} + S \leqslant p < 1, \\ 1 & \text{if } p \geqslant 1. \end{cases} \tag{C.2.6}$$

The shape of the distribution function as well as the expected profit depends on which end of the support firms plays with a positive probability.

If $F(\underline{p}) = 0$. Firms do not put mass at the lower bound of the support and the distribution function is given by:

$$F(p) = \begin{cases} 0 & \text{if } p < V - S, \\ 1 - \frac{V}{p+S} & \text{if } V - S \leqslant p < 1 - S, \\ 1 - V & \text{if } 1 - S < p < V, \\ 2 - \frac{V}{p-S} & \text{if } V \leqslant p < 1, \\ 1 & \text{if } p \geqslant 1, \end{cases} \tag{C.2.7}$$

We proceed by construction. First, we define the lower bound and show that for our value of switching costs, the distribution function must have a mass at the upper bound of the support. Finally, the expression of the expected profits is obtained by the absence of atom at price $\underline{p} + S$.

When firms do not play with positive probability at the lower-bound of the support, then we obtain $\underline{p} = V - S$. If we assume that there is no atom at the upper bound, the expected profit is equal to $\overline{V} = 1 - S$. However, in that case, we obtain that the atom at price $\underline{p} + S$ is equal to $A(\underline{p} + S) = \frac{1-2S(2-S)}{1-2S}$, and this is negative for some of the values of switching costs — a contradiction.

Therefore, the distribution function must have an atom at the upper bound of the support which is $A(1) = \frac{S-1+V}{1-S}$. Furthermore, when setting a price equal to $\underline{p} + S$ no firm attracts nor loses any consumer.[2] Hence, no firm should play such a price with positive probability. The expected profitability is then obtained from expression $A(\underline{p} + S) = 0$. This implies that:

$$A(\underline{p} + S) = 0 \iff 1 - V = 2 - \frac{V}{V - S} \iff V = \frac{S}{V - S}$$

Observe that by our tie breaking rule, a firm setting a price $1 - S$ never loses its captive consumer and attracts the consumer of the rival when the rival sets a price equal to 1. This happens with probability equal to the atom $A(1)$. Therefore, by setting this price, the firm obtains profits:

$$\pi_i(1 - S, \sigma_j) = (1 - S) \times (1 + A(1)) \iff (1 - S)\left(1 + \frac{S - 1 + V}{1 - S}\right) = V$$

Finally, we need to prove that no firm wants to deviate by setting a price out of the support. Clearly, no firm deviates by setting a price above the upper bound of the support as it makes no sales.

if $F(1) = 0$. If firms do not put mass at the upper bound of the support, the distribution function is given by expression

$$\sigma^I = F(p) = \begin{cases} 0 & \text{if } p < \frac{1-S}{2-S}, \\ 1 - \frac{1-S}{p+S} & \text{if } \frac{1-S}{2-S} \leq p < 1 - S, \\ S & \text{if } 1 - S \leq p < \frac{1+S(1-S)}{2-S}, \\ 2 - \frac{1-S}{p-S} & \text{if } \frac{1+S(1-S)}{2-S} \leq p < 1, \\ 1 & \text{if } p \geq 1. \end{cases} \tag{C.2.8}$$

[2] This is the case as the range of the support is twice as low as the value of switching costs.

Again we proceed by construction. First, we obtain the expression of the expected profit by the assumption of no atom at the upper bound. Using the same reasoning as before, there should be a mass at the lower bound. Finally, the expression for the lower bound is obtained by the absence of atom at the price $\underline{p} + S$.

When firms do not play with positive probability at the upper-bound of the support, we obtain $V = 1 - S$. We reach the same contradiction as before by assuming that there is no atom at the lower bound of the support. Therefore, the distribution function must have an atom at the lower bound of the support and it is given by $A(\underline{p}) = \frac{\underline{p}+2S-1}{\underline{p}+S}$. Furthermore, as there is a positive mass at the lower bound of the support, no firm sets with positive probability the price $\underline{p} + S$. Then, the lower bound of the support is obtained from expression $A(\underline{p} + S) = 0$

$$A(\underline{p} + S) = 0 \iff 1 - V = 2 - \frac{V}{\underline{p}} \iff \underline{p} = \frac{1 - S}{2 - S}$$

Observe that by our tie-breaking rule, a firm setting a price equal to $\underline{p} + S$ never attracts the consumer of the rival and it loses its consumer when the other firm sets a price of $\underline{p}$, which happens with probability $A(\underline{p})$. Therefore, by setting this price, the firm obtains profits equal to:

$$\pi_i(\underline{p} + S, \sigma_j) = (\underline{p} + S) \times \left(1 - A(\underline{p})\right) \iff (\underline{p} + S) \times \left(1 - \frac{\underline{p} + 2S - 1}{\underline{p} + S}\right) = 1 - S$$

Proof of Proposition (10) .

Part I: We first prove the case of full collusion. Point (1) of the proposition is the standard case where consumers do not have any switching cost, and we refer to any textbook of Microeconomic Theory for a formal proof. To prove points (2),(3),(4) we define first the present discounted profit of the initial path with full collusion. Because consumers do not switch in the initial path, firms obtain the same profit at each period. Therefore, the present discounted profits are equal to:

$$\Pi_i(\mathbf{p}(0)) = \frac{1}{(1 - \delta)}.$$

If we introduce this result in expression (3.6.1) together with the optimal deviation price p^d and the present discounted profit of the punishment path obtained in proposition (3), we obtain that for a level of switching costs $S \in \left(0, \frac{1}{2+\sqrt{2}}\right]$, the initial path constitutes an equilibrium if:

$$\Pi_i(\mathbf{p}(0)) = \frac{1}{1 - \delta} \geq 2 \times (1 - S) + \delta \left(\frac{(1 + \sqrt{2})S}{(1 - \delta)}\right) = p_i^d + \delta \mathbb{E}(\Pi_i(\mathbf{p}(N)))$$
$$\iff \delta \geq \frac{1 - 2S}{2 - (3 + \sqrt{2})S} = \bar{\delta}(S)$$

For a level of switching costs $S \in \left(\frac{1}{2+\sqrt{2}}, \frac{1}{2}\right)$, whether the initial path is sustained or not depends on the punishment path firms coordinate on. If there is coordination on the mixed strategy equilibrium σ^I, the initial path is sustained if:

$$\begin{aligned}\Pi_i(\mathbf{p}(0)) = \frac{1}{1-\delta} &\geq 2 \times (1-S) + \delta\left(\frac{1-S}{1-\delta}\right) = p_i^d + \delta \mathbb{E}(\Pi_i(\mathbf{p}(N))) \\ \iff \delta &\geq \frac{1-2S}{1-S} = \bar{\delta}(S),\end{aligned}$$

and with coordination on σ^{II}, the initial path is sustained if:

$$\begin{aligned}\Pi_i(\mathbf{p}(0)) = \frac{1}{1-\delta} &\geq 2 \times (1-S) + \frac{\delta V}{1-\delta} = p_i^d + \delta \mathbb{E}(\Pi_i(\mathbf{p}(N))) \\ \iff \delta &\geq \frac{2(1-2S)}{4(1-S) - (S + \sqrt{S(4+S)})} = \bar{\delta}(S)\end{aligned}$$

Finally, point (4) is trivial as the punishment path coincides with the initial path and full collusion is always sustained for any value of the discount factor.

Part II: Now we consider the region where a partial collusive equilibrium is possible. When $S \geq 1/2$, we have already shown that $p^c = 1$ is sustainable. Therefore, we only consider there the case where the magnitude of switching costs is $S < 1/2$. For a collusive price to be sustainable, we need it to be such that no profitable deviation by the firm occurs (condition 1). We also need the collusive price to be higher than the expected profit in the static Nash reversion (condition 2). Because our Nash reversion is one with mixed strategy equilibrium we assume that firms in our model are risk neutral.[3]

We proceed to derive both conditions. Condition (2) is easily computed:

$$p^c \geq (1-\delta) \times \mathbb{E}\left(\Pi\left(\mathbf{p}(N)\right)\right) = V \qquad (2)$$

Condition (1) will depend on the level of switching costs. If $S \leq \frac{p^c}{2}$ we have that the optimal deviation price attracts the consumers of the rival. Therefore, no firm deviates form the collusive price p^c if:

$$\begin{aligned}&\frac{p^c}{1-\delta} \geq 2 \times (p^c - S) + \delta \times \mathbb{E}\left(\Pi\left(\mathbf{p}(N)\right)\right) \\ &\iff (2\delta - 1)p^c \geq \delta(1-\delta) \times \mathbb{E}\left(\Pi\left(\mathbf{p}(N)\right)\right) - 2(1-\delta)S \qquad (1) \\ &(2\delta - 1)p^c \geq \delta V - 2(1-\delta)S\end{aligned}$$

When $S \in \left(\frac{p^c}{2}, \frac{1}{2}\right)$ the optimal deviation price is one arbitrarily close to p^c and the firm will not attract the consumers of the rival, therefore we obtain that the firm does not deviate form the collusive outcome if:

$$\frac{p^c}{1-\delta} \geq p^c + \delta \times \mathbb{E}\left(\Pi\left(\mathbf{p}(N)\right)\right) \iff p^c \geq (1-\delta) \times \mathbb{E}\left(\Pi\left(\mathbf{p}(N)\right)\right) = V \qquad (1)$$

and this coincides with condition (2).

[3] This assumption is crucial to be able to compare profits of the collusive path with Nash.

We just need to verify that both conditions (1) and (2) are satisfied for different values of switching costs. Indeed, in the main text we found the region where a full collusion outcome is sustainable, in this section we look whether partial collusion can be an equilibrium of the game.

$S \in \left(0, \frac{1}{2+\sqrt{2}}\right)$. In this case both equilibria σ^I and σ^{II} have the same discounted profit $V = (1+\sqrt{2})S$. By introducing it in condition (1) we get:

$$p^c(2\delta - 1) \geq S\left[(3+\sqrt{2})\delta - 2\right] \qquad (1)$$

$$\Leftrightarrow \begin{cases} p^c \geq \frac{S[(3+\sqrt{2})\delta-2]}{2\delta-1} \ (=\bar{p}) & \text{for } \delta > 1/2 \\ p^c \leq \bar{p} & \text{for } \delta < 1/2 \\ S = 0 & \text{for } \delta = 1/2 \end{cases}$$

and condition (2) is:

$$p^c \geq (1+\sqrt{2})S$$

- For $\delta > 1/2$: if $\bar{p} > 1$, any form of collusion is impossible, firms win the expected profit earned by playing the Nash reversion. $\bar{p} \leq 1$ implies:

$$\frac{S\left[(3+\sqrt{2})\delta - 2\right]}{2\delta - 1} \leq 1 \iff \delta(S(3+\sqrt{2}) - 2) \leq 2\delta - 1$$
$$\iff \delta \geq \frac{2S - 1}{S(3+\sqrt{2}) - 2} = f(S)$$

  and in the area where there is no full collusion in equilibrium we find that $\bar{p} > 1$. Observe that the last inequality comes from the fact that $S < \frac{2}{3+\sqrt{2}}$ which is the case for the value of switching costs that we are considering here.

- For $\delta < 1/2$: $\bar{p}$ satisfies condition (2) if $\bar{p} > (1+\sqrt{2})S \Leftrightarrow \delta > 1$. So the maximum collusive price in that case is always below the expected gain in Nash reversion and condition (2) is violated.

As a result, for this range of switching costs there is no possible outcome with partial collusion:

$S \in \left(\frac{1}{2+\sqrt{2}}, \frac{1}{2}\right)$ For the equilibrium σ^{II} we know that $V = \frac{S+\sqrt{S(4+S)}}{2}$. By introducing this to condition (1) we get:

$$p^c(2\delta - 1) \geq \frac{5\delta S - 4S + \delta\sqrt{S(4+S)}}{2} \qquad (1)$$

$$\Leftrightarrow \begin{cases} p^c \geq \frac{5\delta S - 4S + \delta\sqrt{S(4+S)}}{2(2\delta-1)} \ (=\bar{p}) & \text{for } \delta > 1/2 \\ p^c \leq \bar{p} & \text{for } \delta < 1/2 \\ 0 \geq \frac{5\delta S - 4S + \delta\sqrt{S(4+S)}}{2} \Leftrightarrow S = 0 \text{ or } S \geq \frac{1}{2} & \text{for } \delta = 1/2 \end{cases}$$

and condition (2) is:

$$p^c \geq \frac{S + \sqrt{S(4+S)}}{2}$$

- For $\delta > 1/2$: if $\bar{p} > 1$, any form of collusion is impossible, firms win the expected profit earned by playing the Nash reversion. $\bar{p} \leq 1$ implies that:

$$\frac{5\delta S - 4S + \delta\sqrt{S(4+S)}}{2(2\delta - 1)} < 1 \iff 5\delta S - 4S + \delta\sqrt{S(4+S)} < 2(2\delta - 1)$$
$$\iff \delta \geq \frac{4S - 2}{5S - 4 + \sqrt{S(4+S)}} = f(S)$$

  and the area where there is no full collusion in equilibrium we find that $\bar{p} > 1$.

- $\delta < 1/2$: $\bar{p}$ satisfies condition (2) if $\Leftrightarrow \bar{p} \geq V$

$$\frac{5\delta S - 4S + \delta\sqrt{S(4+S)}}{2(2\delta - 1)} < \frac{S + \sqrt{S(4+S)}}{2}$$
$$\iff -3S + \sqrt{S(4+S)} \leq \delta \times \left(-3S + \sqrt{S(4+S)}\right) \iff \delta \geq 1$$

  and the profits obtained by Nash reversion are higher than the ones of maximal collusion. Observe that $\sqrt{S(4+S)} > 3S \iff S < \frac{1}{2}$ which is the case that we are considering. We find that, in the equilibrium σ^{II} there exist no equilibrium with partial collusion.

For the equilibrium σ^I we know that $V = 1 - S$. By introducing this to condition (1) we get:

$$p^c(2\delta - 1) \geq \delta - S[2 - \delta] \qquad (1)$$
$$\Leftrightarrow \begin{cases} p^c \geq \frac{\delta - S[2-\delta]}{2\delta - 1} \; (= \bar{p}) & \text{for } \delta > 1/2 \\ p^c \leq \bar{p} & \text{for } \delta < 1/2 \\ S \geq 1/3 & \text{for } \delta = 1/2 \end{cases}$$

and condition (2) is:

$$p^c \geq 1 - S$$

- $\delta > 1/2$: if $\bar{p} > 1$, any form of collusion is impossible, firms earn the expected profit earned by playing the Nash reversion. $\bar{p} \leq 1$ implies that:

$$\frac{\delta - S[2-\delta]}{2\delta - 1} \leq 1 \iff 1 - 2S \leq (1 - S)\delta \iff \delta \geq \frac{1 - 2S}{1 - S} = f(S)$$

  and the area where there is no full collusion in equilibrium we find that $\bar{p} > 1$.

- $\delta < 1/2$: Full collusion is possible if $\bar{p} \geq 1 \Leftrightarrow \delta \geq f(S)$ with $f(S) = \frac{1-2S}{1-S}$. $f(1/3) = 1/2$ and $f(1/2) = 0$. When $\delta < f(S)$, full collusion is not possible but partial partial collusion is possible if $p^c < 1$ is. To check that, we need to verify that $\bar{p}$, i.e. the maximum collusive price, satisfies condition (2): $\bar{p} \geq 1 - S$

$$\frac{\delta - S\left[2-\delta\right]}{2\delta - 1} \geq (1-S) \iff \delta(3S-1) \leq (3S-1)$$

and the above expression is $\delta \geq 1$ if $S \leq \frac{1}{3}$ and $\delta \leq 1$ if $S \geq \frac{1}{3}$. Therefore, for $S \geq 1/3$, the collusion price is higher than the expected profit in Nash reversion players play the highest collusive price possible $p^c = \bar{p} = \frac{\delta - S(2-\delta)}{2\delta - 1}$. For $S < 1/3$, collusion is not possible and players play the Nash reversion.

C.3 Increasing switching costs

Proof of proposition (4) : We start by obtaining a candidate equilibrium in the sub-game where one of the firms have both consumers. Later, we check that no firm wants to deviate. Therefore, we work with the sub-game where one of the firms have both consumers, by deviating at the first period, and for $t-1$ periods there has not been any switch by any consumer. Therefore, at period t there are two types of consumers in the economy. The high switching costs consumer with switching costs S_t and the low switching costs consumer with switching costs S_{t-1}. We define by $\Pi(2) = \sum_{t\geq 0} \delta^t\,(2 \times p_t)$ and $\Pi(1) = \sum_{t\geq 0} \delta^t \times p_t$ the expected discounted profits of having two or one consumer respectively and p_t stand for the equilibrium price. The firm having no consumer obtains zero discounted profits and we use this condition to obtain the equilibrium prices.

When consumers differ on the level of switching costs two possibilities arise. The firm having no consumers can undercut the price to attract both consumers, or it can set a price to attract only the consumer with low switching costs. We proceed to examine both possibilities.

Attract both consumers : The firm having no consumers has zero discounted profits, by undercutting the price by S_t it attracts both consumers. In the continuation game as there is no further switch it gets a discounted profit of $\delta\Pi(2)$. Therefore, the equilibrium price is obtained with the expression:

$$2 \times (p_t - S_t) + \delta\Pi(2) = 0, \tag{C.3.1}$$

and this is equivalent to:

$$2 \times (p_t - S_t) \times \delta^t = -\delta\Pi(2) \times \delta^t \iff \sum_{t\geq 0} (2 \times (p_t - S_t) \times \delta^t) = \sum_{t\geq 0} (-\delta\Pi(2) \times \delta^t)$$
$$\iff \Pi(2) - 2 \times V(S) = -\frac{\delta}{1-\delta} \times \Pi(2) \iff \Pi(2) = 2(1-\delta)V(S).$$

Using this result and equation (C.3.1), we obtain that the punishment equilibrium price is $p_t = S_t - \delta(1-\delta)V(S)$, and the pair of prices or punishment path is equal to:

$$\mathbf{p}(i) = \left\{S_t - \delta(1-\delta)V(S), -\delta(1-\delta)V(S)\right\}.$$

However, with this punishment path, the low switching costs consumer switches to the firm having no consumers. Hence, if the deviating firm wants to ensure that no consumer switches, the difference in prices has to be equal to the magnitude of the low switching cost consumer. Therefore, a punishment path where no consumer switches is given expression (3.7.1) in the Proposition.

Attract low switching costs consumer : If the firm with no consumers aims at attracting only the low switching costs consumers, she undercuts the price of the rival by S_{t-1}. The continuation game that follows, and assuming that there is no further switch, is one where each firm has a consumer each and the discounted profits are given by $\delta\Pi(1)$. Therefore, the equilibrium price is obtained by the expression:

$$(p_t - S_{t-1}) + \delta \times \Pi(1) = 0 \tag{C.3.2}$$

and by doing similar algebra as before, we obtain:

$$(p_t - S_{t-1}) \times \delta^t = -\delta\Pi(1) \times \delta^t \iff \sum_{t\geq 0} ((p_t - S_{t-1}) \times \delta^t) = \sum_{t\geq 0} (-\delta\Pi(1) \times \delta^t)$$

$$\iff \Pi(1) - \delta \times V(S) = -\frac{\delta}{1-\delta} \times \Pi(1) \iff \Pi(1) = \delta(1-\delta)V(S).$$

Introducing this expression into (C.3.2), the equilibrium price is equal to $p_t = S_{t-1} - \delta^2(1-\delta)V(S)$ and the punishment path arising is defined by:

$$p_t(i) = \{S_{t-1} - \delta^2(1-\delta)V(S), -\delta^2(1-\delta)V(S)\},$$

and no consumer switches as the difference between prices is equal to the magnitude of the low switching costs consumers. But then, the way we have obtained this candidate equilibrium price is false as the candidate price equilibrium p_t is different from the one in expression $\Pi(1)$ that emerges when a consumer have switched.

Therefore, the only equilibrium in the incumbency game is:

$$(p_t(\xi = 2), p_t(\xi = 0)) = (S_{t-1} - \delta(1-\delta)V(S), -\delta(1-\delta)V(S)).$$

We now have to check that no firm wants to deviate with the proposed incumbency price. Because consumers are short-sighted, no consumer switches as the price difference is equal to the value of low switching costs consumers. We start by considering possible deviations by the firm with no consumers.

$p_{-i} = -\delta(1-\delta)V(S) - \varepsilon$. There are two possible situations:

$\varepsilon = S_t - S_{t-1}$. This way the firm attracts both consumers and its expected discounted profits are:

$$2 \times (S_{t-1} - S_t - \delta(1-\delta)V(S)) + \delta\Pi(2) \iff 2 \times (S_{t-1} - S_t) - 2\delta(1-\delta)V(S) + 2\delta(1-\delta)V(S)$$
$$\iff 2 \times (S_{t-1} - S_t) < 0 \quad \text{as} \quad S_t > S_{t-1}.$$

$\varepsilon > 0$. This way the firm attracts only the low switching costs consumers and its expected discounted profits are:

$$-\delta(1-\delta)V(S) + \delta\Pi(1) \iff -\delta(1-\delta)V(S) + \delta^2(1-\delta)V(S) \iff -\delta(1-\delta)^2V(S) < 0$$

$p_{-i} = -\delta(1-\delta)V(S) + \varepsilon$. This is not a profitable deviation as the firm does not attract any consumer by setting this price and its expected discounted profits are zero.

Finally, It is left to check that the firm having both consumers does not want to deviate from the proposed punishment path.

$p_i = S_{t-1} - \delta(1-\delta)V(S) - \varepsilon$. It is easy to see that this is not a profitable deviation as the firm has the same demand and it sets a lower price.

$p_i = S_{t-1} - \delta(1-\delta)V(S) + \varepsilon$. For $\varepsilon = S_t - S_{t-1}$. This does not constitute a profitable deviation if:

$$\begin{aligned} 2\times\left(S_{t-1} - \delta(1-\delta)V(S)\right) + \delta\Pi(2) &\geq S_t - \delta(1-\delta)V(S) + \delta\Pi(1) \\ \iff 2\left(S_{t-1} - S_t\right) &\geq \delta\left(\Pi(1) - \Pi(2)\right) + \delta(1-\delta)V(S). \end{aligned}$$

By substituting by the discounted profits that we obtained above and by assuming a comparable value of switching costs $V(S) = (1-\delta)S$ we obtain:

$$\begin{aligned} & 2\left(S_{t-1} - S_t\right) \geq \delta\left(\delta(1-\delta)V(S) - 2(1-\delta)V(S)\right) + \delta(1-\delta)V(S) \\ & \iff 2\left(S_{t-1} - S_t\right) \geq \delta(1-\delta)(\delta-1)V(S) \iff 2\sum_{t\geq 0}\delta^t\left(S_{t-1} - S_t\right) \geq \delta(\delta-1)V(S) \\ & \iff 2S_{t-1} + 2\delta V(S) \geq \left(\delta^2 - \delta + 2\right)V(S) \iff S_0 \geq \left(\frac{\delta^2 - 3\delta + 2}{2(1-\delta)}\right)S. \end{aligned} \tag{C.3.3}$$

Therefore if the initial switching costs is high enough $S_0 \in \left(\left(\frac{\delta^2-3\delta+2}{2(1-\delta)}, 1\right) \times S\right)$. Then, there does not exist a profitable deviation from the proposed incumbency price since we obtain that $\frac{\delta^2-3\delta+2}{2(1-\delta)} < 1 \iff \delta < 1$.

D
Appendix — Chapter 4

D.1 Notes on the Theoretical Model

Our model of revenue management is in the spirit of Talluri and Van Ryzin [51] who study revenue management in the transport industry. But in their model one product consists of a seat and a price. Therefore the same seat proposed at two different prices actually constitutes two different products. In contrast, we consider that the price is a separate variable.[1] Modelling the price as an additional variable implies that the revenue manager never proposes the exact same product at two different prices. This is useful when we allow for some randomness in the preferences of consumers. Indeed in Talluri and Van Ryzin [51], the revenue manager can simultaneously propose the same product at two different prices, and when demand is simulated using a multinomial logit approach, consumers choose the highest price with a positive probability, which is unrealistic.

An additional difference is that we allow for different products to have different capacity constraints. Taking the example of two flights departing the same day, once all the tickets for the first one have been sold, the revenue manager can only propose tickets for the second one.

We show that introducing physically constrained characteristics has an impact on the revenue manager's strategy. Talluri and Van Ryzin [51], show that the revenue manager only proposes menus on the efficient frontier when optimizing her profit function. For them, a menu is efficient if linear combinations of other menus cannot do better in terms of expected revenue $\mathbb{E}(\pi|\tilde{p})$ except by increasing the total probability of purchase. We show that this is not necessarily the case in our model. We take the example of two types of products, each limited to 200 units. Possible prices for each product are $P = \{50, 70, 90, 110\}$. Consumer preferences are defined as in section 4.5. In Figure D.1.1, we plotted all combinations of products and possible prices according to their total probability of purchase and

[1] We can consider that in Talluri and Van Ryzin [51], the revenue manager can only control which types of products to propose. Our revenue manager has control over two dimensions: which types of products to propose and at what price.

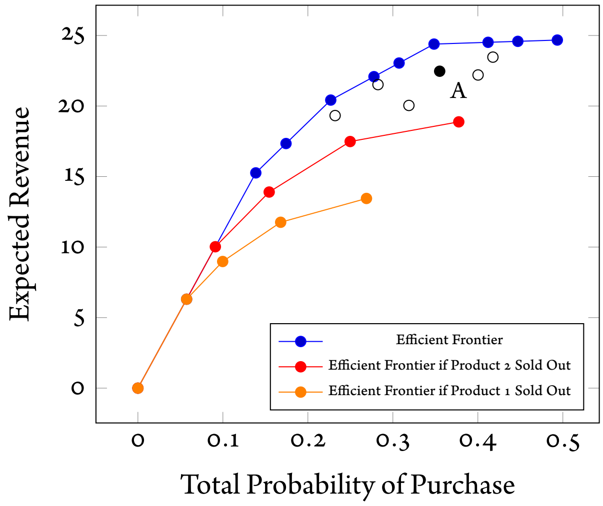

Figure D.1.1: Total probability of purchase (horizontal axis) and the associated expected revenue (vertical axis) of the different possible menus of prices. Efficient frontiers are represented by plain lines. Menu A, which is not on the efficient frontier, is played 12.9% of the time when the arrival rate of the consumer is $\lambda = 0.6$.

expected revenue. Product-price combinations that are efficient according to Talluri and Van Ryzin [51] are represented by coloured dots and linked together. According to the theoretical result of Talluri and Van Ryzin [51], other combinations should not be proposed by the revenue manager. However, by computationally solving for the optimal pricing of the revenue manager, when $\lambda = 0.6$, the revenue manager proposes an inefficient menu of prices, $(90, 50)$, in 12.9% of the periods of the booking process. This percentage corresponds to an average over 1000 simulations. We attribute this to the fact that we allow for several capacity constraints.

D.2 Values of Parameters

In Table D.2.1, we summarize the values we use for consumer preferences and train characteristics. Simulations (1) refer to the simulations in the distributive properties of the revenue management. Simulations (2) refer to the case of heterogeneous consumers. Simulations (3) and (4) respectively stand for indirect and direct competition. Finally, simulations (5) are used in robustness checks presented in Appendix D.3.

Table D.2.1: Summary of parameters' values used in the different simulations.

	(1)	(2)	(3)	(4)	(5)
# trains	1	2	1	2	2
X_i	400	200	400	200	200
ν	1.5	1.5	1.5	1.5	1.5
α	0	-0.5	0	0	-0.5
γ	0.03	$\gamma^B = 0.01$; $\gamma^L = 0.03$	0.03	0.03	0.03
ν_0	1	1	$\{0.1, 0.5, 1, 1.5, 2\}$	1	1
T	2000	2000	2000	2000	2000
# simulations	1000	1000	1000	500	1000

D.3 Discussion of Revenue Management Vs. Optimal Fixed Price

We define the load rate as the ratio between the number of sold units over the total number of available units of the same type. If we take the example of a plane with load rates of 100% for first-class tickets and 80% for economy seats, it means that all seats in first class were booked and 80% of the economy seats in the plane were booked. It is therefore a measure of how full the plane is. Compared to the fixed price strategy, the increase in consumer surplus is driven by higher amounts of sold units or load rates. This is due to the finer choice set of the revenue manager which allows her to post higher or lower prices depending on the past sales. For instance, for $\lambda = 0.9$, when the spread around the optimal price is equal to 5, the revenue manager posts the lowest and highest prices with a relatively high frequency: respectively 21% and 17% of the time. Of course, the larger the spread, the smaller the frequencies at which the lowest and highest prices are posted since they become dominated by the fixed price.

This seems to indicate that revenue management is especially useful for large market sizes (at least when consumers are homogeneous) because in that case it allows the revenue manager to temporarily decrease the posted price when there is no sales and to take some risk in the other case. When demand is low and it is nearly impossible to sell all available units, sales variations never reach the point which induces the revenue manager to change the optimal fixed price. Intertemporal discrimination in this case becomes useless as the changes in the probability of reaching the capacity constraint are negligible compared to the coarseness of the price grid.

Figure D.3.1 displays the average posted price as a function of the booking period for different intensities of demand and a small spread, $\xi = 7$. Obviously, the average level of posted prices increases with the arrival rate of consumers. We also observe that the optimal fixed price is played with probability 1 at the beginning of the booking period. Then, in the middle of the booking period, the average posted price is roughly equal to the optimal fixed price, with some noise or without, respectively when $\lambda = 0.7, 0.9$ and when $\lambda = 0.3$. When $\lambda = 0.5$, it is even slightly higher. This indicates that deviations from the optimal fixed price may occur if the market is sufficiently large, but these deviations might be centred around the optimal fixed price, i.e. the revenue manager might have on average equal incentives to lower or increase the price. At the end of the booking period, the revenue manager tends to post the lowest price more often. This is consistent with the incentives derived from the theoretical model.

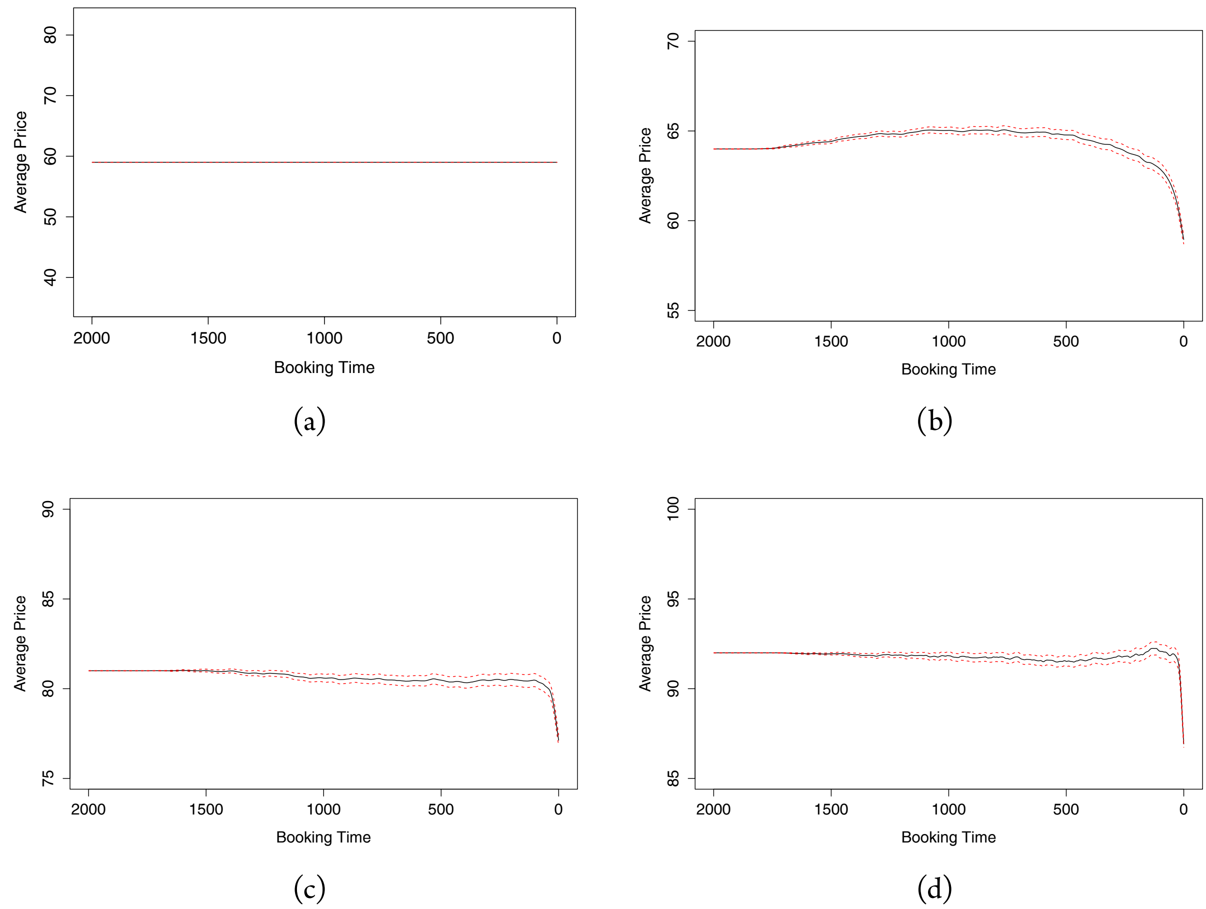

Figure D.3.1: Change in posted average prices as a function of the booking period when the price set is $P = \{p^o - 7, p^o, p^o + 7\}$. Figures (a),(b),(c),(d) deal with increasing intensities of demand, corresponding to an expected arrival of respectively 600, 1000, 1400, and 1800 consumers.

Remark 6. *To test whether this result is robust to a more complete choice set, we do the same exercise allowing the choice set of the revenue manager to be richer. In these simulations:*

$$P = \{p^o, p^o \pm 5i | 1 \leq i \leq \xi\} \;\; \textit{for}\; \xi = 1, 2, \ldots$$

As shown by Figures D.8.4 and D.8.5 of Appendix D.8.2, we find similar qualitative results in that case. Consumer surplus does not decrease even if the action set of the revenue manager is richer. The average price *throughout the booking period is again very close to the average optimal price (within 1 unit) but the load rates are significantly higher when λ is sufficiently high and the revenue manager actually discriminates. For instance, at $\lambda = 0.9$, the load rate under revenue management is around 99.7% against 98% for optimal pricing.*

All these findings are robust to a more general specification including different substitutable types of products.

In a more realistic framework, the revenue manager has to optimize her profit over two types of products from which consumers can choose. We want to test whether or not adding more flexibility in the choices of consumers affects the comparison between revenue management and a fixed price strategy. We suppose that the choice set P of prices available to the revenue manager remains identical between the two types of products. In the case of a fixed price strategy, we also assume that the proposed price is identical for both types. In transports, this could correspond to a situation in which the revenue manager is required to set an identical price for all transports running on a particular origin-destination leg. The parameters used in this section can be found in Table D.2.1 of Appendix D.2, column (5). $a = -0.5$ means that consumers now prefer product 1 over 2. For instance, if the two types of products refer to two different flights the same day on the same origin-destination leg, this can be interpreted as one flight being more convenient than the other.

We test whether or not revenue management is an improvement over a fixed price strategy and measure to which extent revenue management including the optimal fixed price raises profits and affects consumer surplus.

To do so, we compute the optimal fixed price p^o for the two types of products and the simulated average profit and consumer surplus associated with this price. We then construct the choice set P of the revenue manager to include the optimal fixed price and some small variations around it:

$$P = \{p^o - \xi, p^o, p^o + \xi\}_{i=1,2} \; \xi = 1, 2, \ldots$$

The revenue manager's choice set is again small but Figures D.8.6 and D.8.7 of Appendix D.8.2 show that revenue management strongly increases both profits and consumer surplus. In fact, all the conclusions we found for one type of products seem to extend to this case. For small market sizes, the effect is not significant. For $\lambda \geq 0.4$, revenue management increases both profits and consumer surplus. For some spread values, profits rise from 4% for $\lambda = 0.5$ up to 12% for $\lambda = 0.9$. This corresponds to large increases in consumer surplus: between 20% and 30% for the same spread values.

The importance of the impact on profit in this case is easily explained by the way we construct the optimal fixed price: we consider a unique fixed price for two types of products, i.e. rule out any possibility to price discriminate between the two types even though one is more attractive than the other. Introducing revenue management offers this possibility despite identical choice sets for the two types. This gives more leeway to the company and considerably raises profits. However, price discrimination between the two types is only profitable when the arrival rate is sufficiently high.

What is more surprising is the huge positive impact on consumer surplus. One might expect that price discrimination between the two types of products allows the firm to extract consumer surplus more easily. We attribute this result to the increase in load rates: more consumers can afford to purchase a product, especially at the beginning of the booking period. For instance, for $\lambda = 0.5$, 88% of the seats of the less attractive product are sold under revenue management with a spread $\xi = 15$ whereas only 61% of the seats are sold under a fixed price strategy.

Remark 7. *To isolate the sole impact of revenue management as intertemporal price discrimination from the impact of price discrimination between types of products, we also carry out an analysis in which we compute an optimal fixed price for each type. The resulting price vector is denoted* $\mathbf{p^o}$. *We then construct the choice set P of the revenue manager to include for each type, its optimal fixed price and some small variations around this price:*

$$P = \{p_i^o - \xi, p_i^o, p_i^o + \xi\}_{i=1,2} \; \xi = 1, 2, \ldots$$

In this case, the results are similar in nature although the importance of the impact of revenue management is much weaker, as seen in Figures D.8.8 and D.8.9 of Appendix D.8.2. For $\lambda \geq 0.4$, *the increase in profit lies between 0.8% and 1.7%, with a more pronounced impact for high intensities of demand. For the same arrival rates, the increase in consumer surplus varies between 1% and 2.7%.*

D.4 Extension to Heterogeneous Consumers

D.4.1 Increasing Arrival Rates

To simulate two types of consumers randomly arriving at each period, we first have to simulate the probability that one consumer arrives at a given period, then draw the type. Indeed, we cannot have two consumers arriving during the same period. Since in the application to transports, the business type has a probability of arrival increasing as time comes closer to departure and the arrival probability of the leisure type is constant, the overall probability of arrival must increase as well.

Considering this, we define the probability of arrival of one consumer as :$\lambda_t(\rho_l, \rho_{b_t}) = \frac{\rho_l + \rho_b}{\rho_l + 1}$, where ρ_l and ρ_b are the relative rates of arrival of respectively leisure and business passengers. The probabilities of being of type $\tau (= b, l)$ conditional on an arrival are defined as:

$$P(\tau | arrival) = \frac{\rho_\tau}{\rho_\tau + \rho_{-\tau}} \qquad \tau = b, l$$

The unconditional probability of a type-τ arrival is therefore given by:

$$P(\tau) = P(\tau \cap arrival) = \frac{\rho_\tau}{\rho_l + 1} \qquad \tau = b, l$$

We find this way of modelling arrivals convenient since it implies that for all values of ρ_l, $\lambda_t(\rho_l, .)$ goes to 1 as ρ_b goes to 1, which allows us to model a very intense demand towards the end of the booking period.

We want the arrival rate of leisure types to be constant and the one of business types to increase in time. Hence ρ_b has to vary across time. We model it as a discretized version of an exponential distribution:

$$\rho_{b_t} = e^{-\frac{t}{\mu}} \qquad \text{for } t = T, \ldots, 1$$

Here, μ is a large number which we use to parametrize the curvature of ρ_{b_t}. As t is large, i.e. far from the departure date, ρ_{b_t} is close to zero and tends to 1 as t goes to 1.

Table D.4.1 gives the correspondence between the expected number of arrivals we want to model and our two parameters, μ and ρ_l.

Table D.4.1: Correspondence between the expected number of arrivals, μ, and ρ_l.

$\mathbb{E}(arrivals)$	400	600	800	1000	1200	1400	1600	1800
μ	223	355	511	710	988	1430	2290	4815
ρ_l	$\frac{1}{9}$	$\frac{3}{17}$	$\frac{1}{4}$	$\frac{1}{3}$	$\frac{3}{7}$	$\frac{7}{13}$	$\frac{2}{3}$	$\frac{9}{11}$

D.4.2 The Modified Bellman Equation

The Bellman equation of section 4.3 naturally extends to this new set-up:

$$\begin{aligned} V_t(X) = \max_{\tilde{p}\in\Pi_i P_i} \sum_{\tau\in\{b,l\}} \sum_i & Pr_t(\tau)Pr(d_i = 1|\tilde{p},\tau)\,[\tilde{p}_i + V_{t-1}(X_i - 1, X_{-i})] \\ & + [1 - \lambda_t + Pr(l)Pr\,(d_0 = 1|\tilde{p}, \tau = l) \\ & + Pr_t(b)Pr\,(d_0 = 1|\tilde{p}, \tau = b)]\,V_{t-1}(X) \end{aligned} \tag{D.4.1}$$

subject to the final constraints:

$$\begin{cases} V_0(X_i) = 0 & , \forall X_i \\ V_t(X) = 0 & , \text{whenever } \sum_i X_i = 0 \\ \tilde{p}_i = +\infty & , \text{if } X_i = 0 \end{cases} \tag{D.4.2}$$

The demand mechanism is modelled as before except that coefficients in the utility functions are type-dependent.

D.5 A Reduced-Form Model of Indirect Competition

To evaluate the impact of indirect competition on revenue management, we study how average profits, consumer surplus and load rates change with the value of the outside option.

D.5.1 Demand Modelling with Variable Outside Option

The consumer utility when she chooses the outside option is now: $u_0 = \ln(v_0) + \varepsilon_0$, in which $v_0 \in \mathbb{R}^{+*}$ is a parameter capturing the competition intensity and ε_0 is the random part of the utility defined as before. In the benchmark model of section 4.3.3, $v_0 = 1$. The higher v_0, the higher the intensity of indirect competition.

This new specification yields the following market shares:

$$\begin{cases} Pr(d_i = 1|\tilde{p}) = \frac{e^{\bar{u}(p_i)}}{\sum_{j\in I} e^{\bar{u}(\tilde{p}_j)} + e^{\ln(v_0)+\mathbb{E}(\varepsilon_0)}} = \frac{e^{\bar{u}(p_i)}}{\sum_{j\in I} e^{\bar{u}(\tilde{p}_j)} + v_0} \\ Pr(d_0 = 1|\tilde{p}) = \frac{1}{\sum_{j\in I} e^{\bar{u}(\tilde{p}_j)} + e^{\ln(v_0)+\mathbb{E}(\varepsilon_0)}} = \frac{v_0}{\sum_{j\in I} e^{\bar{u}(\tilde{p}_j)} + v_0} \end{cases} \tag{D.5.1}$$

It is obvious that:

$$\lim_{v_0 \to 0} Pr(d_0 = 1|\tilde{p}) = 0 \text{ and } \lim_{v_0 \to \infty} Pr(d_0 = 1|\tilde{p}) = 1$$

The other features of the model remain the same.

D.5.2 Values of Parameters and Results

Values of parameters used in this setup are summarized in Table D.2.1 of Appendix D.2, column (3). The arrival rates are the same as before, i.e. $\lambda \in \{0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9\}$, and the choice set of the revenue manager is given by $P = \{20, 40, 60, 80, 100, 120, 140, 160\}$.

Figures D.8.16 and D.8.17 of Appendix D.8.4 summarize our results for indirect competition: as competition increases, profits go down while consumer surplus increases. Since we model an increase in competition as an increase in the value of the outside option, these results are quite intuitive: a higher outside option means that the revenue manager's products are less attractive (hence less profits) whereas it also mechanically raises the final consumer surplus through the utility of consumers who do not purchase.

We also look at the impact of indirect competition on average prices posted during the booking period and on price dispersion. These results are summarized in Figure D.8.18 of Appendix D.8.4. Unsurprisingly, stronger indirect competition drives prices down. However, its impact on price dispersion is more ambiguous: for low arrival rates, price dispersion is higher when competition is weak; the opposite is true for high arrival rates.

D.6 Mixed Equilibria

In a mixed equilibrium, players randomize over their action set, here: $\tilde{P} = \{\underline{p}^i, \bar{p}^i, +\infty\}$, $i = R, A$. We have shown that in this special case, $\tilde{p}^i = +\infty$ is always dominated by the other prices unless $X^i = 0$. Therefore, any mixed equilibrium of a continuation game at t is a randomization over $\{\underline{p}^i, \bar{p}^i\}$ for $i = R, A$. In the following, we denote $\sigma^i = P(\tilde{p}^i = \underline{p}^i)$ and by abuse of notation σ^i also denotes the action of player i when she randomizes.[2]

Finding an equilibrium in mixed strategies of the continuation game at t amounts to find σ^i for $i = R, A$:

Suppose i plays σ^i. Then, $-i$ needs to be indifferent between $\underline{p}^{-i}$ and $\bar{p}^{-i}$ to play a mixed strategy, which is summarized by the following condition:

$$\varphi^{-i*}_t(X, \sigma^i, \underline{p}^{-i}) = \varphi^{-i*}_t(X, \sigma^i, \bar{p}^{-i}) \tag{D.6.1}$$

where:

[2]Although any pure strategy is a mixed strategy, here we refer to *mixed strategies* if and only if $\sigma^i \in (0, 1)$

$$\varphi^{-i*}_t(X, \sigma^i, p) = \sigma^i \varphi^{-i*}_t(X, \underline{p}^i, p) + (1 - \sigma^i)\varphi^{-i*}_t(X, \bar{p}^i, p)$$

Equation D.6.1 therefore yields:

$$\sigma^i = \frac{\varphi^{-i*}_t(X, \bar{p}^i, \bar{p}^{-i}) - \varphi^{-i*}_t(X, \bar{p}^i, \underline{p}^{-i})}{\varphi^{-i*}_t(X, \bar{p}^i, \bar{p}^{-i}) - \varphi^{-i*}_t(X, \underline{p}^i, \bar{p}^{-i}) + \varphi^{-i*}_t(X, \underline{p}^i, \underline{p}^{-i}) - \varphi^{-i*}_t(X, \bar{p}^i, \underline{p}^{-i})} \tag{D.6.2}$$

D.7 Multinomial logit approach and horizontal differentiation

Here are more detailed explanations about why overall demand for the proposed products increases if we introduce an additional vertically undifferentiated choice. Assume that given a vector of price p, the utility of buying any product is given by u whether we are in the monopolistic or the duopolistic case.

Then, in the monopolistic case, the probability of buying a product is:

$$p^m(buy|p) = \frac{e^u}{e^u + 1} = 1 - \frac{1}{e^u + 1}$$

In the duopolistic case, the probability of taking of choosing the product of one firm (e.g. firm 1) is given by:

$$p^d(buy\,1|p) = \frac{e^u}{2e^u + 1}$$

which is of course lower than the probability of buying a product in the monopoly. However, the overall probability of buying in the duopoly is given by $1 - \frac{1}{2e^u+1}$, which is higher than in the monopoly.

To give some economic intuition to this result, we say that adding an additional type of products, even vertically undifferentiated, can create horizontal differentiation.

D.8 Simulation Results

D.8.1 Dynamics of Prices

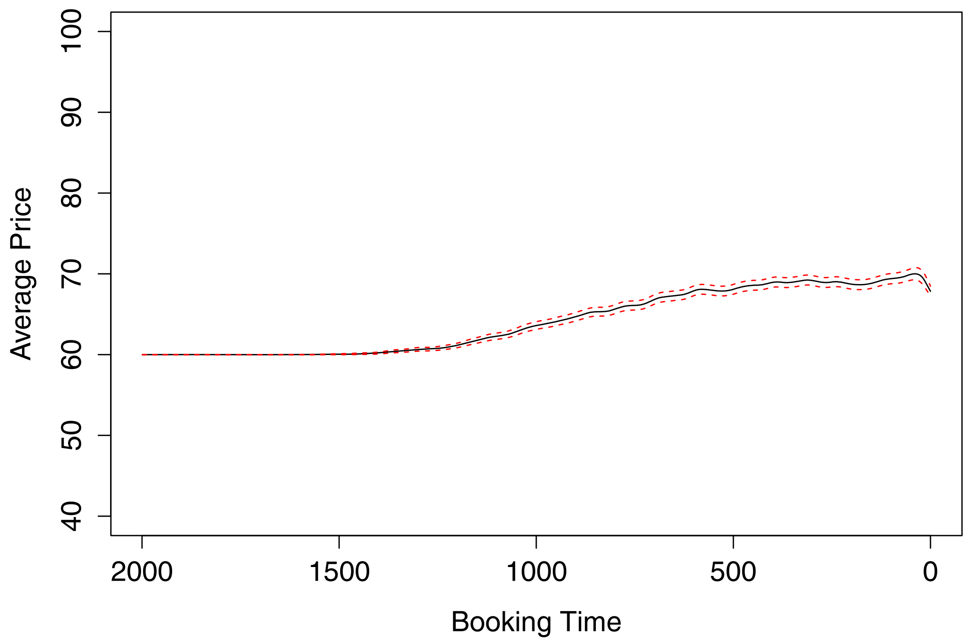

Figure D.8.1: Average change in prices when $\lambda = 0.5$ and the revenue manager chooses the price from $\{20, 40, 60, 80, 100\}$.

D.8.2 The Distributive Properties of Revenue Management

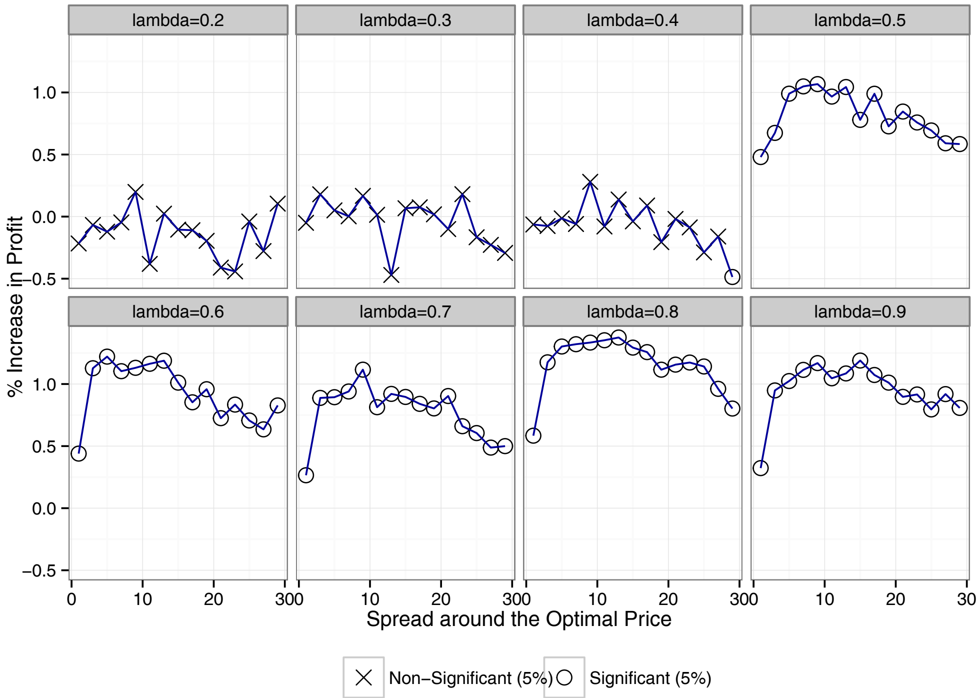

Figure D.8.2: Average change in profit between optimal fixed pricing and revenue management for different intensities of demand for homogeneous products.

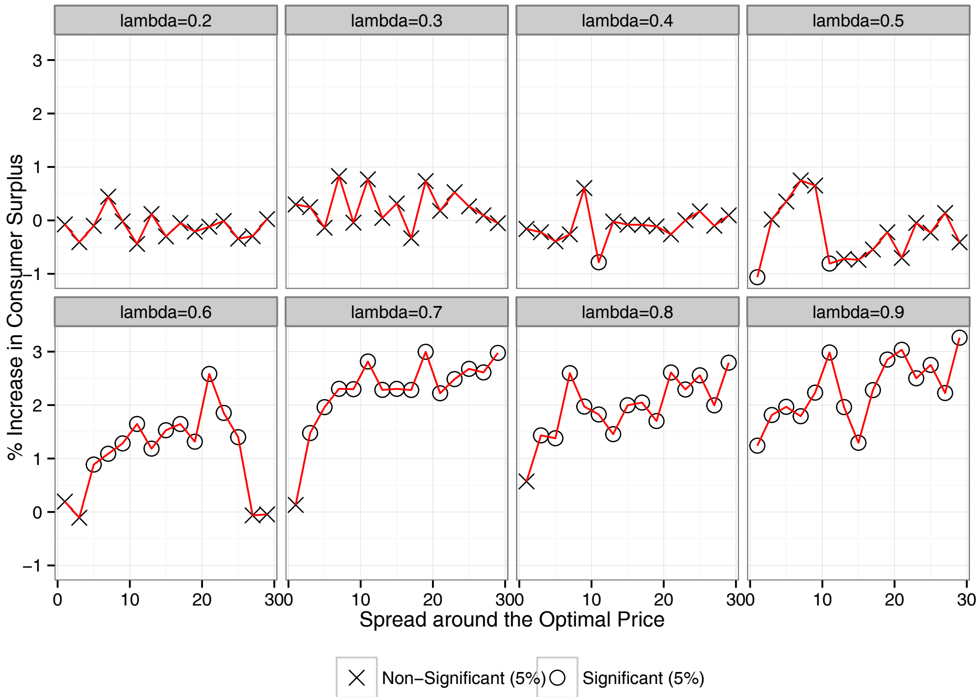

Figure D.8.3: Average change in consumer surplus between optimal fixed pricing and revenue management for different intensities of demand for homogeneous products.

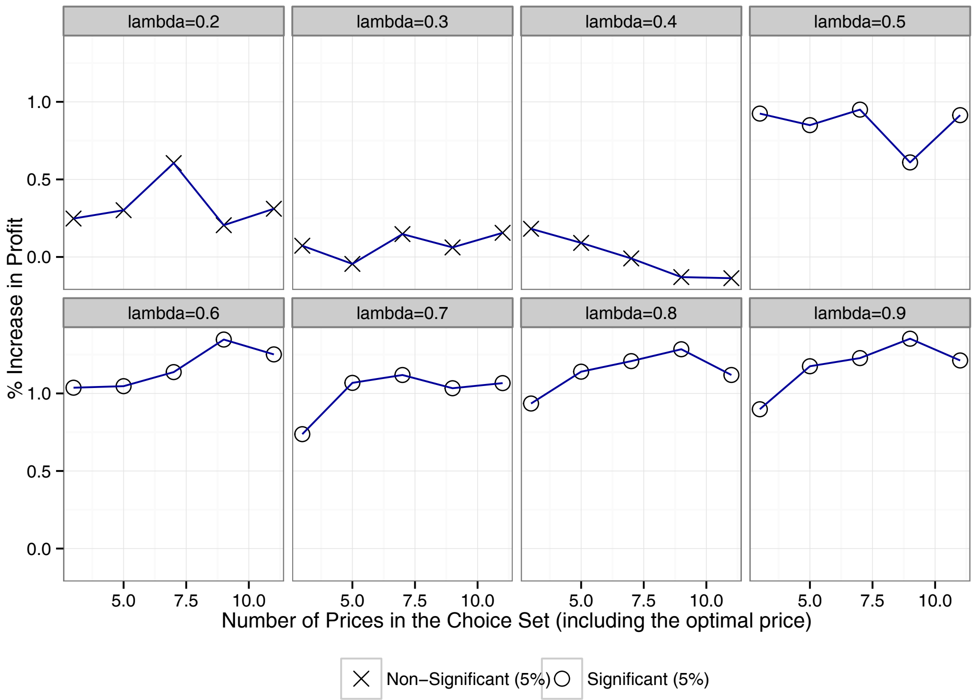

Figure D.8.4: Average change in profit between optimal fixed pricing and revenue management for different intensities of demand for homogeneous products and various sizes of choice sets.

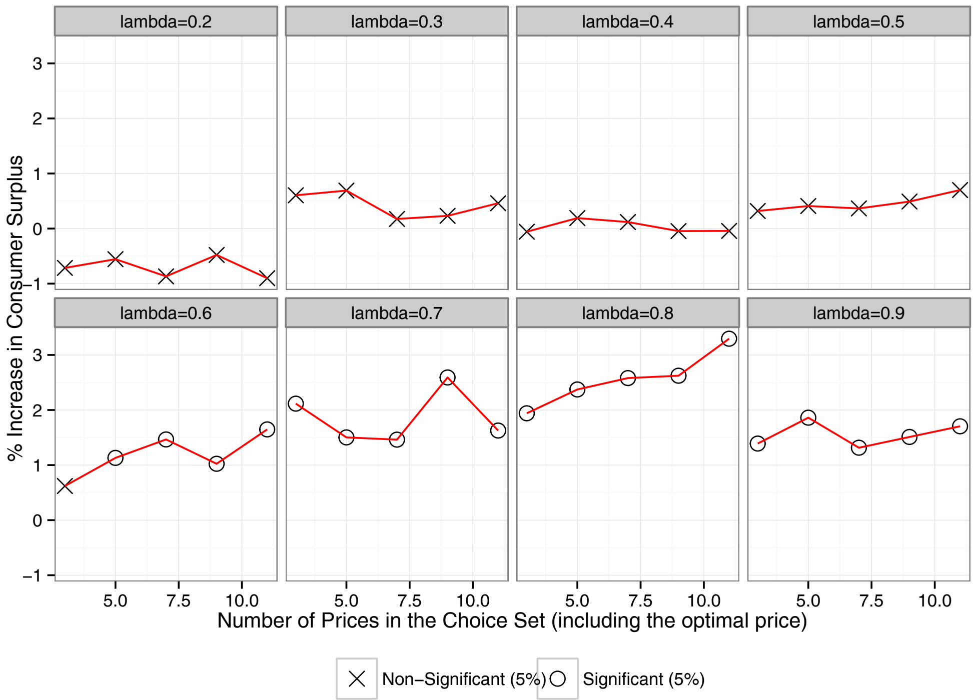

Figure D.8.5: Average change in consumer surplus between optimal fixed pricing and revenue management for different intensities of demand for homogeneous products and various sizes of choice sets.

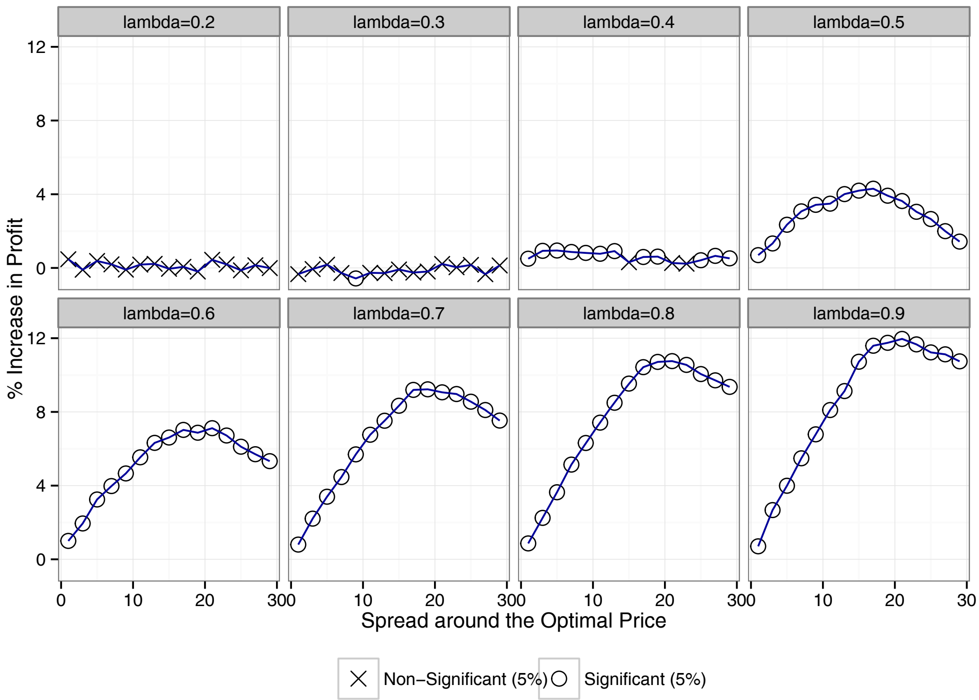

Figure D.8.6: Average change in profit between optimal fixed pricing and revenue management for different intensities of demand and two types of products. The optimal fixed pricing consists here of a unique optimal price for all products.

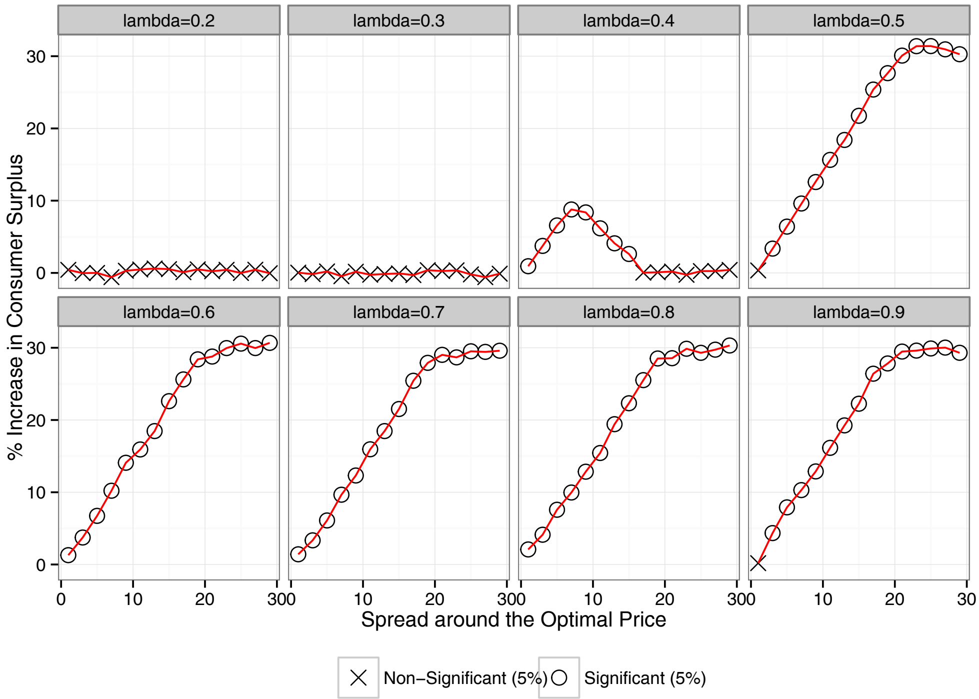

Figure D.8.7: Average change in consumer surplus between optimal fixed pricing and revenue management for different intensities of demand and two types of products. The optimal fixed pricing consists here of a unique optimal price for all products.

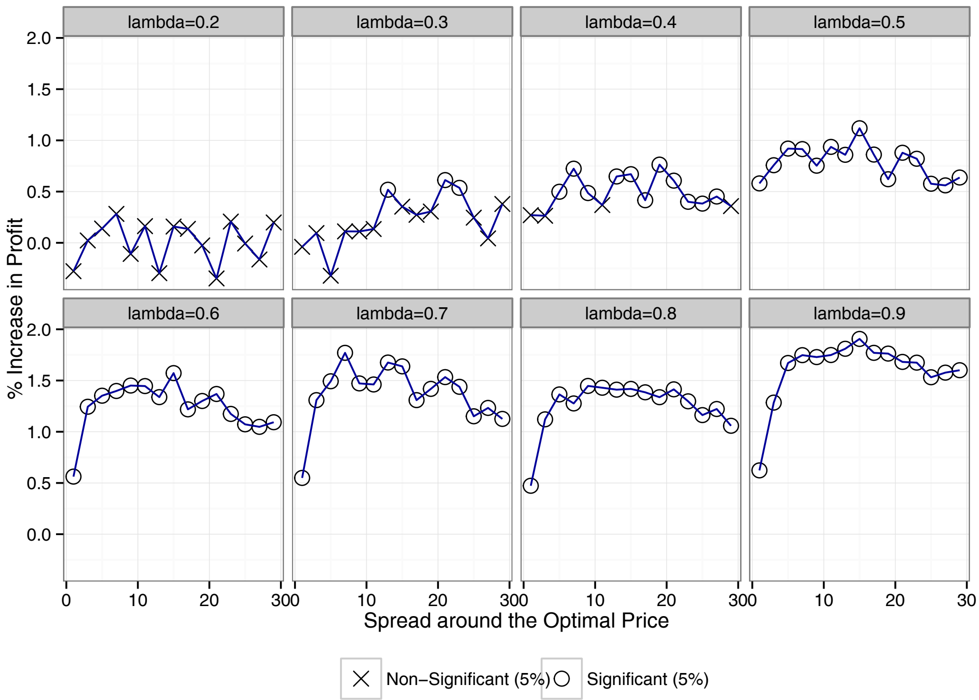

Figure D.8.8: Average change in profit between optimal fixed pricing and revenue management for different intensities of demand and two types of products. We compute here a fixed optimal price for each type of products.

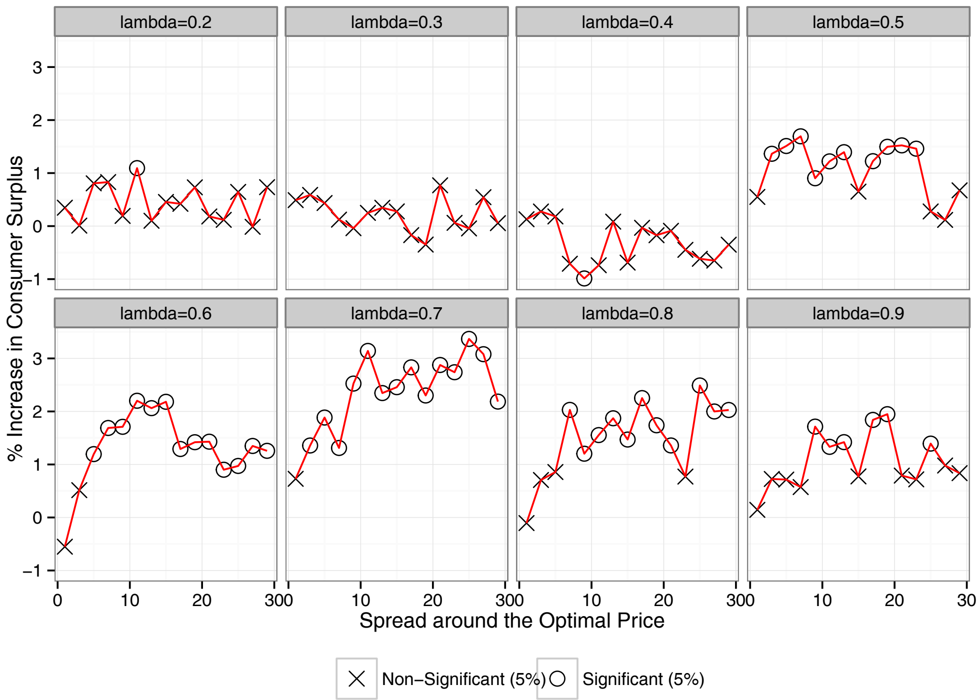

Figure D.8.9: Average change in consumer surplus between optimal fixed pricing and revenue management for different intensities of demand and two types of products. We compute here a fixed optimal price for each type of products.

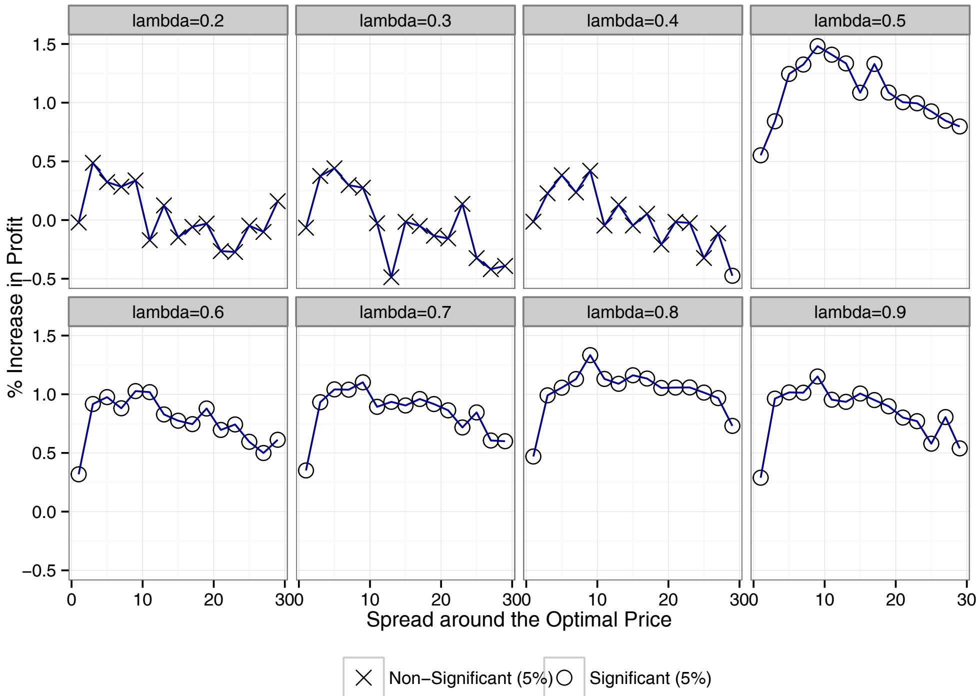

Figure D.8.10: Average change in profit between optimal fixed pricing and revenue management for different intensities of demand when consumers have an increasing willingness to purchase.

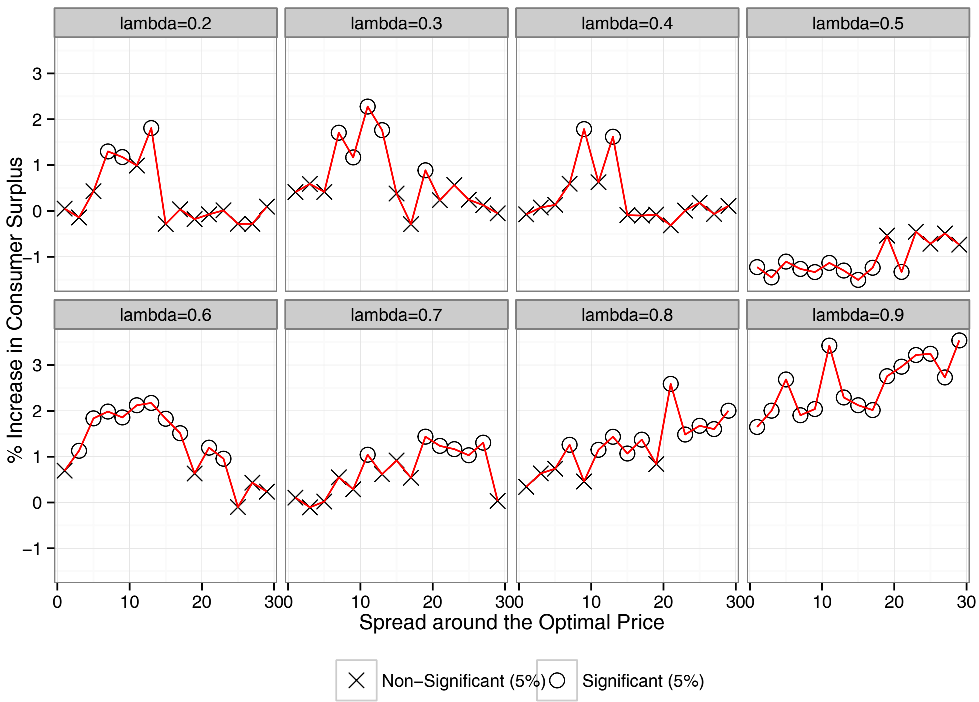

Figure D.8.11: Average change in consumer surplus between optimal fixed pricing and revenue management for different intensities of demand when consumers have an increasing willingness to purchase.

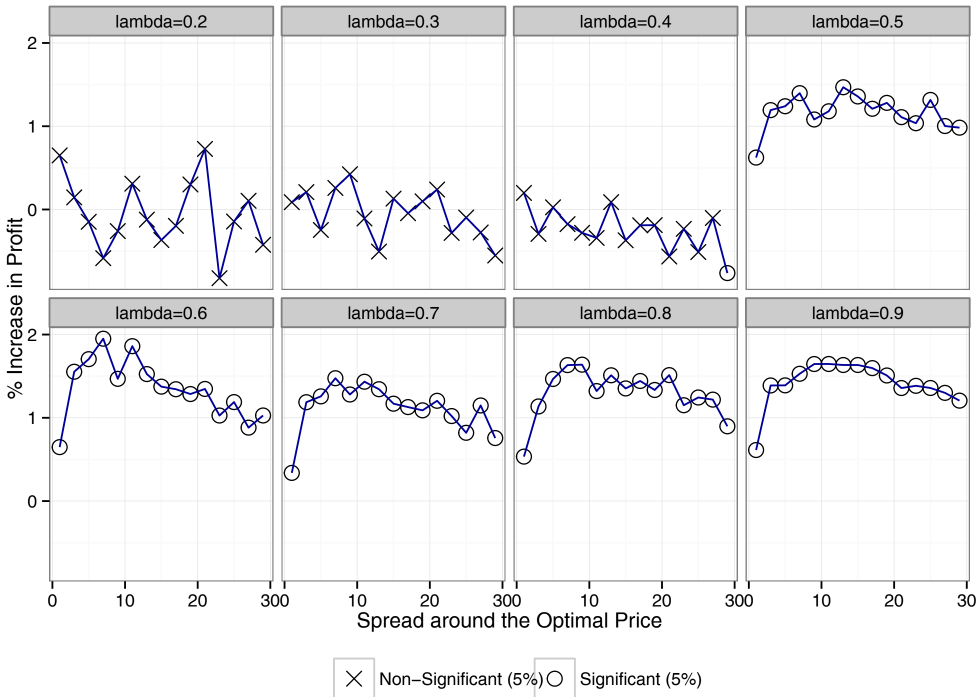

Figure D.8.12: Average change in profit between optimal fixed pricing and revenue management for different intensities of demand when the arrival rates of consumers is noisy.

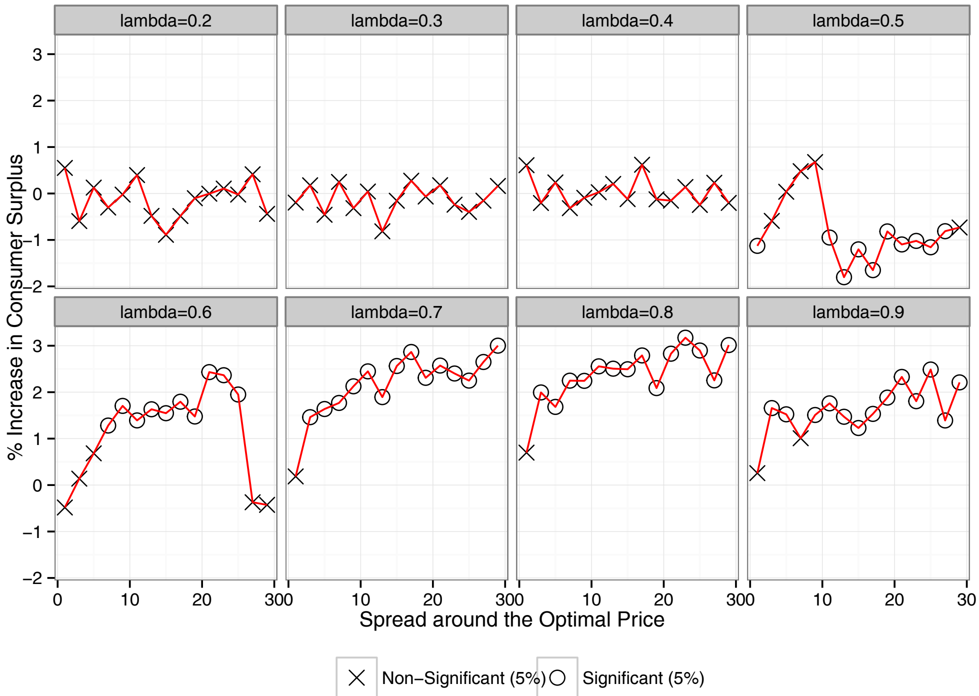

Figure D.8.13: Average change in consumer surplus between optimal fixed pricing and revenue management for different intensities of demand when the arrival rate of consumers is noisy.

D.8.3 Revenue Management and Heterogeneous Consumers

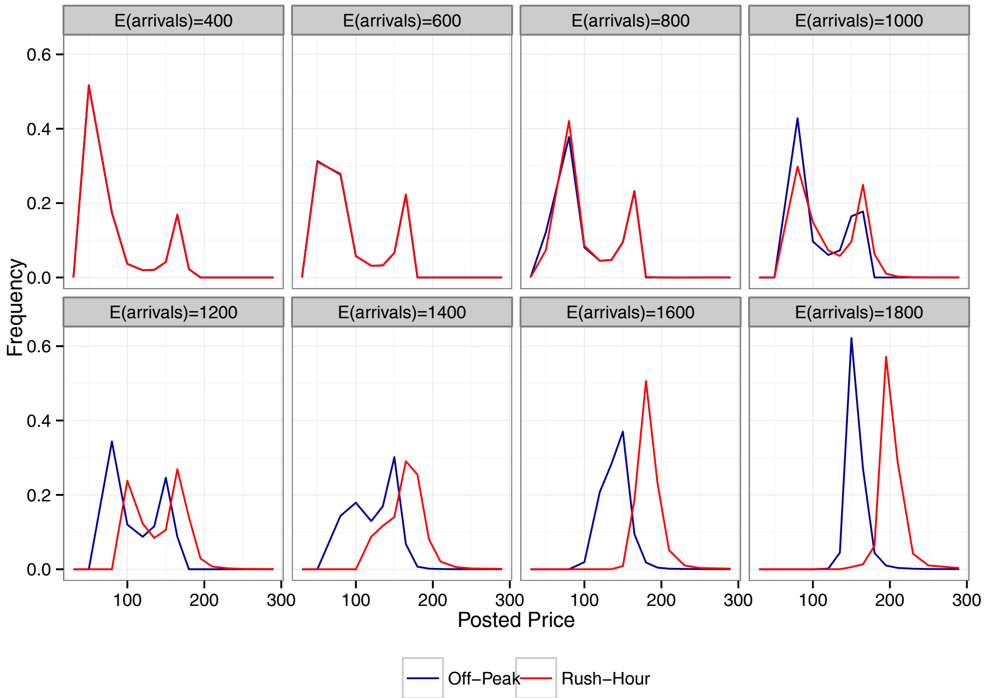

Figure D.8.14: Distribution of the revenue manager's choices for each type of products in the case of heterogeneous consumers.

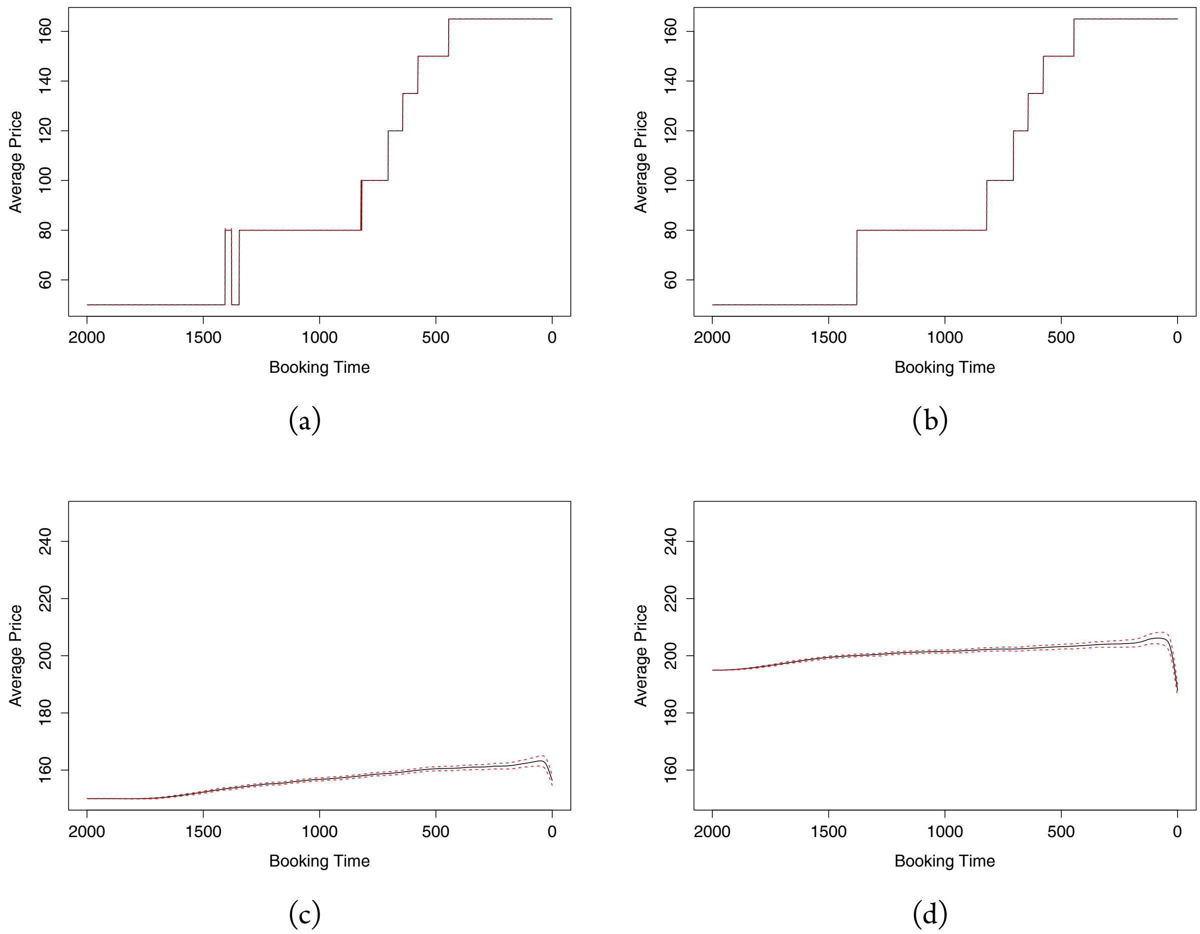

Figure D.8.15: Change in posted average prices as a function of the booking period in the case of heterogeneous consumers. Figures (a) and (b) respectively deal with the off-peak and rush-hour trains for an expected arrival of 600 consumers. Figures (c) and (d) respectively deal with the off-peak and rush-hour trains for an expected arrival of 1800 consumers.

D.8.4 Revenue Management and Intermodal Competition

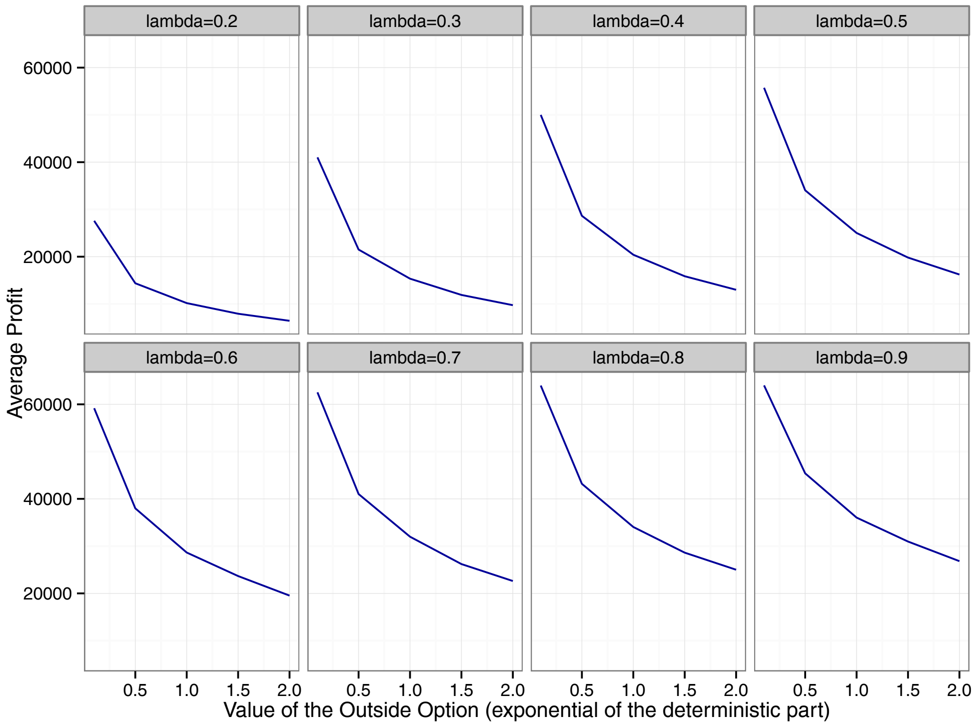

Figure D.8.16: Average profit generated by revenue management for different market sizes and different values of the outside option.

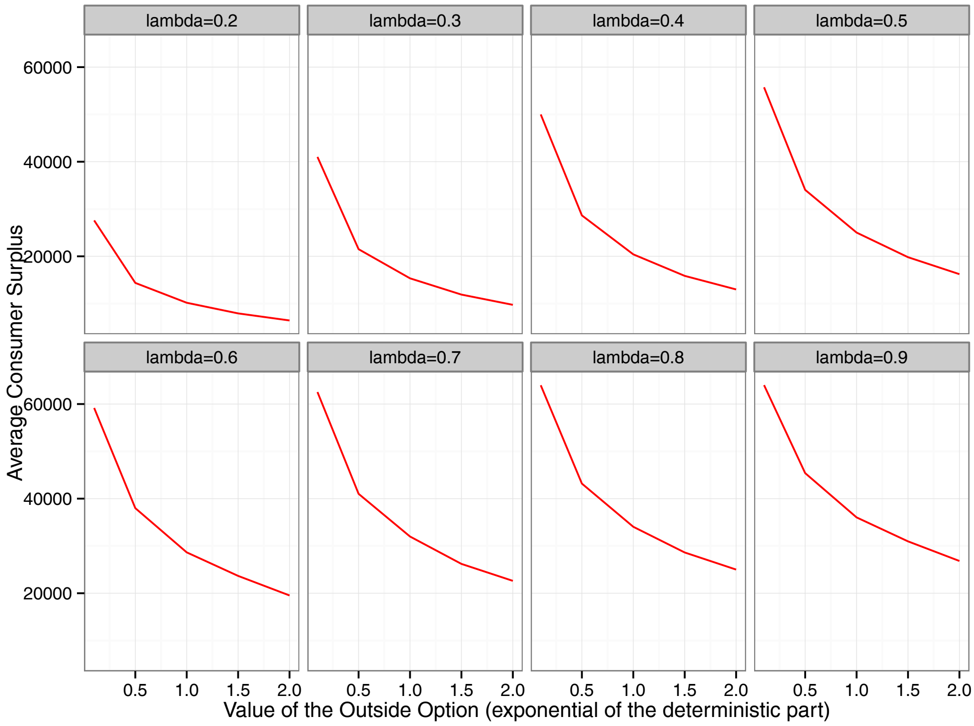

Figure D.8.17: Average consumer surplus generated by revenue management for different market sizes and different values of the outside option.

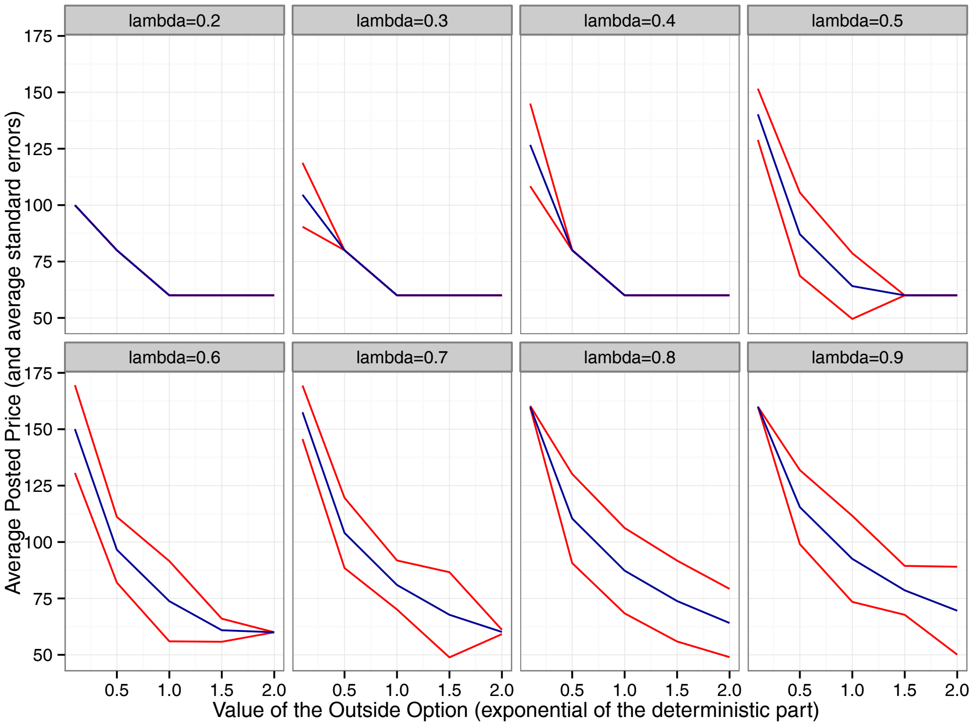

Figure D.8.18: Average price and price dispersion in the case of revenue management for different market sizes and different values of the outside option.